AmCham EU
SPEAKING FOR AMERICAN BUSINESS IN EUROPE

Consultation response

Artificial intelligence in the financial sector

AmCham EU speaks for American companies committed to Europe on trade, investment and competitiveness issues. It aims to ensure a growth-orientated business and investment climate in Europe. AmCham EU facilitates the resolution of transatlantic issues that impact business and plays a role in creating better understanding of EU and US positions on business matters. Aggregate US investment in Europe totalled more than €3.7 trillion in 2022, directly supports more than 4.9 million jobs in Europe, and generates billions of euros annually in income, trade and research and development.

American Chamber of Commerce to the European Union
Speaking for American business in Europe

Avenue des Arts/Kunstlaan 56, 1000 Brussels, Belgium • **T** +32 2 513 68 92
info@amchameu.eu • amchameu.eu • European Transparency Register: 5265780509-97

Executive summary

Artificial intelligence (AI) has been used by financial institutions for many years to enhance efficiencies, improve risk management and foster innovation. At the same time, like other technologies, AI may introduce challenges for firms in relation to data governance, regulatory compliance and risk oversight. Therefore, a regulatory framework that balances innovation and risk management while aligning with global standards is necessary. To achieve this, policymakers should maintain a principles-based and technology-neutral regulatory approach that builds on existing frameworks to avoid unnecessary complexity.

While the EU AI Act provides an important step towards harmonising AI governance, some areas would benefit from further clarification through technical standards or delegated acts, such as the definition of AI. However, since financial institutions are already familiar with the technology and managing the associated risks in line with existing financial services regulations, sector-specific guidance for financial services is not needed at this stage. Any potential future guidance should emerge from a collaborative approach among financial institutions, third-party AI providers and regulatory bodies, taking into account existing financial services legislation and risk management practices to ensure proportionality and no additional regulatory complexity.

Introduction

The European Commission has launched a targeted consultation on AI in the finance sector, seeking to collect information on the application and impact of AI in financial services. The consultation aims to understand the opportunities, challenges and risks associated with AI's growing role in financial services and to evaluate how regulatory frameworks can support the safe and responsible use of these transformative technologies. In response to this consultation, the American Chamber of Commerce to the EU (AmCham EU) provides insights on behalf of the industry on the following topics included in the consultation:

- How AI has been used in the financial sector to date;
- The implications of the use of generative AI in the financial sector;
- How financial institutions are managing AI-related risks;
- Future challenges in scaling AI;
- The probability of future systemic risks; and
- The applicability of the EU AI Act to financial services.

1. AI use cases

The use of AI is not new to financial services, and some financial institutions have been exploring and adopting the use of AI and machine learning (ML) applications for more than a decade. Generally, use cases in production across financial services firms today contribute, amongst other things, to:

- Enhancing back office functions such as risk management and regulatory compliance through facilitating the aggregation and consolidation of information.
- Improving outcomes for clients (for example, improving the response time and quality of responses to queries and more targeted products and offers).
- Enhancing employee productivity and improving operational efficiency by automating manual processing tasks. Examples include using optical character recognition – the process that converts an image of text into a machine-readable text format – and natural language processing (NLP) – machine learning technology that gives computers the ability to interpret, manipulate and comprehend human language.

AI technology can be separated into 'traditional' AI, meaning the use of AI techniques primarily focused on data analysis and making classifications, predictions or decisions based on that data, including the use of ML; and generative AI, which is capable of generating new content in various forms based on probabilistic assessments and can be adapted to a wide range of tasks with minimal training.[1]

1.1 Traditional AI

Banks in particular have for several years been leveraging traditional AI technology across different lines of business. Examples of use cases in banks today include:

- Fraud detection: Machine learning models to detect fraud, for instance on outbound payments and inbound deposits.
- Trading: Reinforcement learning techniques to improve pricing and hedging accuracy. For example, an AI tool may generate bid/offer recommendations for certain markets or products when a trader receives a request. The trader can then choose whether to accept the recommendation or modify it.
- Anti-money laundering/transaction monitoring and sanctions screening: AI tools for high volume businesses (like payments) to screen transactions and detect anomalies in transactions. AI/ML models may also be used to reduce false positives and improve operational efficiency.
- Credit decisioning: Traditional ML models based on decision trees are used to facilitate credit decisions about credit approvals, detecting overlimit account transactions, pricing and loan amounts for customers. ML models can also be used to underwrite risk scores.
- Marketing and customer support: Banks are using AI tools to engage and retain existing customers by using historical customer data to identify next best products and target rewards offers. NLP has also been used to help direct and resolve customer queries.
- Cybersecurity: Banks are starting to use AI to detect and respond to potential cyberattacks more quickly and efficiently. For instance, AI can be used by security analysts to help classify suspicious emails. However, this should be combined with human verification.

In addition to the above, firms are also exploring traditional AI in back office functions, including financial reporting, knowledge management and employee productivity enhancements.

[1] Newer models are emerging that can ingest and understand multimodal inputs, such as the combination of audio, video and text information. Currently, these emerging variants of AI are only gradually being adopted in the financial services context and largely in the context of low-risk applications, subject to risk management frameworks.

1.2 Generative AI

1.2.1. Generative AI technology

The ability of generative AI technology and large language models (LLMs) to understand and generate natural language and to unlock informational value from unstructured data has huge value potential for firms. These emerging techniques can help firms to transform data and analytics for faster, better-informed decision-making, improve operational efficiencies and better manage risk, regulations and fraud.

Nonetheless, generative AI models have inherent limitations and give rise to certain unique challenges and risks. These include susceptibility to hallucinations, limited explainability of the AI model, firms' reliance on a limited number of third-party vendors and data privacy risks. As a consequence, firms are exploring the use of generative AI cautiously, and applications are generally low risk, in early stage use and with a human-in-the-loop. Where necessary, financial institutions are adapting their existing risk management frameworks to account for the risks presented by generative AI, and third-party providers of AI solutions are proactively benchmarking themselves against established and developing AI governance frameworks (see further detail on risk management frameworks under section 2.1).

1.2.2. Generative AI use cases

As noted, firms are exploring generative AI in a controlled manner, with most use cases currently in development being lower risk and internal facing, targeted at improving productivity. Examples of early use cases include:

- Research: Accelerating research and increasing the quality of decision-making by searching across proprietary datasets, structured and unstructured fields and websites.
- Strategic analysis: Developing mergers and acquisitions (M&A) theses and assisting financial statement analysis.
- Data analysis & summarisation: Providing ad-hoc portfolio exposure information and summarising attribution for portfolios.
- Collaboration: Improving productivity by helping locate and summarise information and improve communication between internal groups and customers.
- Fraud: Enhancing fraud detection and querying large amounts of information at speed and scale.
- Risk mitigation: Interacting with trade facilitators to better understand settlement failures and how to resolve them.
- Regulatory change management monitoring: Translating changes in regulatory and business requirements into code.
- Compliance: Detecting anomalies, automating manual controls and overseeing new product delivery compliance.
- Software development: Reducing the amount of time it takes for developers to write, translate or debug code by using natural language descriptions.

Some firms are also experimenting with customer-facing generative AI, use cases including chatbots for customer interaction and in marketing and sales, through the creation of customised client content. Generative AI may also be used to better analyse customer behaviour to develop personalised product offerings.

1.3 Benefits

Specific benefits derived from the use of AI will be dependent on the use case and the business models of firms deploying those use cases. However, as outlined in the previous section, there are some overarching benefits of both traditional and generative AI for firms and clients.

1.3.1. Benefits for firms

Most current benefits for individual firms resulting from the use AI/ML are derived from improving efficiencies through automation and other techniques to optimise manual processes, in addition to using AI/ML for risk management and regulatory compliance functions. Increasing these efficiencies provides cost savings benefits and allows for a more effective allocation of resources, freeing up employees to focus on more complex and higher-value tasks. Further benefits include enhanced fraud detection (thereby reducing fraud-related financial losses for firms), enhanced decision-making and improved cybersecurity operations through contextualised and actionable visibility into the latest security threats. The use of AI/ML in credit can also benefit firms by facilitating better risk management through more informed credit decisions.

1.3.2. Benefits for clients

The use of AI/ML can also deliver better outcomes for customers. The use of AI in credit can allow financial institutions to better assess credit risk. All else being equal, a more predictive credit scoring model can allow financial institutions to expand access to credit for customers. AI can also be used to tailor products and services to clients and provide enhanced customer support, for example through the use of NLP in call centres to help identify why customers are calling and provide intelligent call routing services. AI/ML methods in fraud to detect suspicious activity also contribute to a reduction in fraud rate for customers and a better experience, as more legitimate transactions are approved at the point-of-sale.

1.4 Development of AI models: in-house vs third-party vendors

Whether firms develop in-house models or rely on third-party vendors will depend on a variety of factors, including the specific use case and the resources available to the institution. Financial institutions are generally exploring both in-house and third-party applications.

Cloud service providers provide computing power, infrastructure and access to AI tools to financial institutions of all sizes, allowing them to leverage their unique data and information streamed into large-scale, real-time databases and apply AI models to perform high-value functions, including the assessment, quantification and calculation of financial risk.

Firms may opt to develop in-house solutions to leverage proprietary datasets or where third-party vendor models aren't available for a specific use case. This requires firms to have sufficient resource to be able to develop models in house. For firms facing resourcing obstacles, the use of open-source data could be an option, although this would not always be suitable where data is proprietary.

A lack of available resources and an absence of large datasets mean theoretically smaller firms may be less able to develop internal models than larger firms. These smaller financial institutions may be more sensitive to regulatory uncertainty or lack the necessary time and resources to interpret regulatory expectations. Therefore, continuous and clear expression of regulator receptivity to the adoption of new technologies is crucial to ensure equal footing for these institutions.

Yet, hedge funds, including smaller ones, may also opt to develop internal models (if they have the necessary resources) based on their own datasets to support investment decisions for proprietary reasons, rather than use third-party vendors. Equally, it might be simpler and more cost-effective for banks to use third-party vendors, including for non-finance specific, process-driven applications (eg summarisation) even if they have the resource to develop their own models. For generative AI models, firms are generally using third-party solutions rather than developing in-house solutions.

2. Risk management

Many of the risks presented by AI applications are not novel or specific to the use of AI and are risks that firms are already expert in managing through existing risk management frameworks. Some risks, however, could be amplified by the use of AI, including data privacy, bias and discrimination, transparency and explainability, and cybersecurity risks. To manage and address these risks, firms are taking a disciplined approach to risk management and leveraging existing firm-wide risk management frameworks.

As financial regulators have made clear, consumer protection and financial market laws and regulations apply to financial activities involving the use of AI just as they would to the use of any other technology. Where AI amplifies existing risks or presents new risks and as the technology evolves, institutions will continue to uplift their frameworks accordingly and abide by comprehensive AI governance principles as well as guidance by policymakers on specific AI risks. Risk management frameworks help financial institutions and third parties ensure that they satisfy regulatory expectations. To this end, the broad, principles-based approach to model risk management guidance continues to provide an appropriate framework for managing AI model risk.

Ultimately, banks have robust and mature risk management frameworks in place to meet numerous existing regulatory and supervisory requirements relating to technology and are therefore well-equipped to manage risks from AI.

2.1 Risk management frameworks

In assessing risk, AI models are not inherently riskier than non-AI models. A risk-tiering assessment must consider the context or activity for which a model is used, as well as the model's complexity and materiality. To assist in these assessments, regulators could clarify that the use of AI or generative AI alone does not place a model into a high-risk tier and publish further cross-sectoral guidance to help set expectations regarding the materiality/risk ratings of AI models as applied to common use cases. It's also worth highlighting that the specific risks and mitigations depend on the business use of generative AI and the related data sensitivity implications, and that the use of AI may not be appropriate for all use cases.

2.1.1. Model risk governance

Firms are already subject to comprehensive regulatory guidance on model risk in the EU, including the European Central Bank's guide to internal models,[2] and have in place well-established model risk governance frameworks to meet regulatory requirements. These frameworks ensure oversight of the model development process, including testing, assessing conceptual soundness of models, confirming underlying data, considering model complexity and transparency, assessing and evaluating

[2] https://www.bankingsupervision.europa.eu/legalframework/publiccons/pdf/ssm.pubcon230622_guide.en.pdf

implementation and on-going performance. This framework applies to all models used within bank processes, irrespective of the underlying technology used.

Given the unique characteristics of AI technologies, some aspects of how risk management frameworks apply to AI models may be less clear, as noted above. This uncertainty may impede industry's progress developing and adopting AI-based models. Absent clarity in these circumstances, firms may be incentivised to take overly conservative approaches that could result in longer lead times to production or in initiatives being deprioritised without a clear path to implementation.

2.1.2. Technology controls

Regulators set standards and expectations for sound risk management and evaluate controls for the use of technology, including AI technologies, which firms are responsible for meeting. Sound technology risk management by firms include maintaining an inventory of AI technologies being implemented, assessing the level of risk associated with each use case, expectations for testing and on-going validation, and issue and incident tracking. Effective information security, cybersecurity, resilience, privacy, and operational and fraud-related controls are also important for the use of AI.

2.1.3. Third-party risk management

Firms manage third-party AI risks through existing third-party oversight (TPO) programmes, policies and processes that establish an overarching and technology-neutral third-party risk management (TPRM) framework. TPRM frameworks are designed to be able to adapt to changes in technology and business models, including those stemming from emerging technologies such as AI and generative AI. Many of the types of third-party risks that AI presents are not new or unique to AI and are therefore adequately managed by the protections embedded throughout the lifecycle of a firm's existing TPRM program (ie pre-onboarding risk assessments, due diligence processes, supplier control requirements and existing contractual frameworks).

Many of the risks associated with third-party use of AI can be traced back to the data used to train and run AI models, rather than the AI systems or algorithms themselves. Whilst not specific to AI, the concerns are around: the reliability and quality of the data (which can lead to hallucinations); legal issues around protection of such data; the potential for biases to be found within datasets and to enter the training process; the potential for reputational risk from unauthorised use of data or where generative AI models are introduced into an existing data use case that was previously authorised.

Transparency, a primary challenge for AI, is also relevant in the TPRM context. However, it is important to distinguish between model risk and third-party risk when using third-party AI models. Transparency in a model risk context concerns the clarity and interpretability of the model's inner workings and decision-making processes. Where appropriate, a developer of an underlying foundation model should provide documentation outlining how the model is intended to be used, known inappropriate uses, known risks and recommendations for deployers and users to manage risk. Without this, implementing robust model validation processes and testing procedures in respect to third-party models can be challenging. Additionally, developers and deployers should implement rigorous, tailored best practices such as: addressing safety and harm prior to deployment; implementing systematic internal reviews grounded in policies and guiding principles; employing a high bar for evaluations; sharing and leveraging AI responsibility tools; and continuing to advance mitigations. Careful thought and consideration should be applied in situations where misuse might occur.

Transparency in the context of TPRM, on the other hand, relates to the ability for a bank to manage the risks of the arrangement, including the existing or planned use of AI, banks' ability to perform due diligence on a third-party's control environment, their compliance with regulations and adherence to contractual obligations. In this context, the ability for third-parties to explain why, when and how AI is being used, information about data inputs and outputs and with what governance and risk management measures is important. This challenge extends to awareness of potential vulnerabilities

and compliance issues across the supply chain to ensure firms are appropriately protected from the risks that AI introduces.

Adapting TPRM frameworks to address any AI-specific challenges would involve taking a risk-based approach and adjusting or adapting control requirements as needed. Examples include conducting due diligence at onboarding and throughout the third-party arrangement to identify where and how AI is being used, or renegotiating third-party contractual arrangements to ensure third-parties are obligated to notify firms of their use or planned use of data in connection with an AI model or capability.

As the use of AI by third parties expands, and advancements and innovations in AI continue to introduce new challenges and risks, it is increasingly important to be proactive about understanding how third parties are using or planning to use AI and emerging AI technologies. This highlights the importance of flexible, technology-neutral and outcomes-focused regulatory frameworks to manage third-party risk, including those related to AI. This enables TPO programmes to evolve, ensuring robust risk management that keeps pace with technological advancements, whilst also managing the associated risks. Importantly, it also allows banks to structure their approach to AI risk management and address the interdisciplinary nature of AI risk beyond TPRM frameworks and across various functions such as cybersecurity, technology, AI governance teams, business resiliency, model risk management and operational risk management.

2.1.4. Data governance

Effective data governance for AI tools is critical. Firms process various information including non-personal, personal and financial information every day, and they have been investing for several years in enhancing data governance frameworks to reduce data-related risks and to comply with existing requirements on managing data risk. Furthermore, the General Data Protection Regulation (GDPR) obligates firms to ensure data is accurate across its lifecycle. Organisations adopting International Organization for Standardization (ISO) certification have incorporated the required data governance requirements as a baseline standard within their data governance programmes. As part of their data governance frameworks, firms generally have implemented centrally developed data management policies, internal standards and related committees. These can be applied to manage risks relating to the use of AI tools and developed to effectively safeguard against further data risks resulting from the use of AI (such as those relating to data quality, accuracy, data movement, information security and privacy). Priorities for firms in progressing data governance standards for AI include working to enhance data quality by better identifying and remediating data quality issues, monitoring data use and controls to govern which data can be used for specific use cases.

Generative AI models pose particularly complex data challenges due to the large datasets on which such models tend to be trained as well as the ability for such models to create new information based on that data. Whilst firms can adapt existing data governance frameworks to account for these risks (and are doing so where they are deploying LLMs), data standardisation could help promote data interoperability and address certain data-related concerns with respect to training and implementing AI models and tools.

2.1.5. Generative AI governance

Generative AI is another technique used for modelling for which governance frameworks need to be adapted, rather than needing entirely new frameworks. As with other AI risks, firms are adapting existing AI governance structures to accommodate generative AI. These structures are accountable for overseeing the internal deployment of generative AI use cases. Most use cases have a human-in-the-loop responsible for the use of the output. These users must have the expertise and tools to verify and/or challenge the outputs. To mitigate the risk of hallucinations, firms employ methods to improve response accuracy, eg using prompt engineering, whereby instructions are provided to the model

designed to achieve the desired and accurate response and/or using Retrieval Augmented Generation stores to combine authoritative external knowledge bases with training data to optimise outputs.

2.2 Challenges to scaling

Although financial services firms are managing AI-related risks through existing governance and risk management frameworks, they face challenges when it comes to scaling AI models, particularly generative AI models. Data, explainability and access to talent are key obstacles to scaling.

2.2.1. Data

The availability of high-quality, diverse and representative data that is organised and accessible is crucial for firms to be able to scale models, particularly LLMs. Meeting the appropriate standard of data quality requires a significant amount of preparatory work by firms. Although they should be held accountable for the quality of data they prepare for AI models, it is important firms are not held to an impossible standard of data quality. Any requirements relating to data quality should account for the fact that training, validation and testing data sets can never be completely free of errors or bias. Privacy protections are typically embedded into AI systems by design, ensuring that data collection, processing and storage are aligned with regulatory requirements. The application of privacy principles should guide the data collection, pre-training and fine-tuning of LLMs. Organisations that develop or deploy generative AI models should be transparent and accountable for explaining the privacy principles they follow and maintaining an internal privacy programme that documents their privacy practices.

From a regulatory perspective, privacy and data protection rules are becoming increasingly intertwined in global economic competition given the importance of data in the growing technology economy. That is leading to a somewhat fragmented or divergent approach to these rules in some cases and a lack of harmonisation may lead to data portability barriers between jurisdictions. For example, GDPR regulators' application of the GDPR in relation to AI systems is still developing, leading to inconsistent positions on for example, the appropriate legal basis for training AI systems and even whether AI systems contain personal data at all, creating an uncertain legal landscape.

To address this, the Commission should continue to encourage the use of consistency mechanisms through the European Data Protection Board to ensure the consistent application of GDPR across Member States. Likewise, in its ongoing work on promoting cross-border payments the Financial Stability Board (FSB) has recently identified friction between national data frameworks as a barrier to interoperability.[3] Similar barriers may extend to other areas of technology like AI, for which cross-border data sharing is crucial to ensuring optimal functionality. On a firm-specific basis, regulators should keep in mind that firms may have less of an incentive to develop datasets if it is likely those datasets will be subject to forced sharing, which ultimately may hamper innovation. It is important to

[3] https://www.fsb.org/wp-content/uploads/P16072-1.pdf

consider live data sharing initiatives (such as the European Commission's Open Finance/Financial Data Access proposal) in the context of innovation.

2.2.2. Explainability

Explainability refers to explaining why a model generated a given output. Depending on the particular use case or application, the required degree of – and approach to – explainability may vary. Additionally, the ability to trace back and explain outcomes from AI systems operating at scale may differ depending on the type of AI used.

Complex AI models, particularly generative AI models, are often less transparent because their decision-making processes are not easily understandable by humans. However, just because some AI models are less explainable (compared to non-AI models), it should not always be assumed these models are a 'black box' that cannot be interpreted. Due to the dynamic nature of generative AI models and the different options available, reliance on extensive and ongoing testing focused on outcomes throughout the development and implementation stages of such models should often be prioritised relative to explainability in satisfying regulatory expectations of soundness. To that end, the development of technical metrics and related testing benchmarks should be encouraged. Model 'explainability', while useful for understanding the specific outputs of AI models, may be less effective or insufficient for establishing whether the model as a whole is sound and fit for purpose. Critically, as new model types become available so will the technical solutions for explainable AI. Advances in explainable AI, including developments in explainability diagnostic techniques, are helping firms better understand models and overcoming explainability limitations. Nevertheless, outcome-driven assessments and continuous monitoring, validation and stress testing will remain essential to maintain AI model reliability.

Firms should endeavour to build models with the appropriate levels of explainability across use cases. From a governance and risk management perspective, it would not be appropriate to have uniform explainability requirements for different use cases; expectations of explainability should be based on the context in which an AI model is being used and the recipient of the explanation. For example, the explainability of a system is more important for some customer-facing applications (eg credit decisioning) compared to internal applications used by firms and employees of firms for simple processing tasks. It is important that developers and deployers understand regulator expectations of documentation requirements for all categories of AI risk management. Sufficiency of documentation should be determined by what is needed for firms to use and validate the model and understand its design, theory and logic.

2.2.3. Talent

Developing, implementing and maintaining AI models requires experts in AI, data science and machine learning who also understand the specific requirements and challenges of the financial services sector. That same level of expertise is also necessary in the public sector. Although banks have been hiring new employees with AI-specific skills, there is a shortage of professionals with the requisite skill to work in the financial sector, which can make it difficult to scale AI initiatives.

2.3 Systemic risk

Whilst firms are well equipped to manage micro-financial risks at a firm level, there is increasing concern about the systemic risks arising from the widespread use of AI in the financial sector, including those relating to concentration risk and algorithmic trading. The assessment of these risks is explained below, but as banks are generally proceeding slowly and cautiously with AI applications, the use of AI in finance does not currently present any systemic risks. However, these risks are worth regular monitoring as the use of AI progresses.

2.3.1. Concentration risk

It is recognised that the growing use of technologies within financial services has introduced questions around the extent to which this use might lead to concentrated risk exposures on a limited number of suppliers in certain areas. Both firms and regulators have been navigating and addressing the associated operational and cybersecurity risks at both a firm and system level across the broader spectrum of technology services, including through frameworks such as the Digital Operational Resilience Act (DORA). As the use of AI and generative AI rises, and AI becomes more integrated into various operational frameworks, the potential for concentration risk will need to be managed. In this context, existing risk management frameworks continue to be appropriate for identifying and managing the risks stemming from potential supplier concentrations (as outlined in section 2.1.3). Although we do not believe the concentration of LLM providers is presently a financial stability risk, these concerns should continue to be monitored as AI take-up grows.

2.3.2. Algorithmic coordination

There are concerns that the use of the same or similar models by market participants might result in the risk of herding whereby algorithmic outputs lead to uniformity in behaviour and potentially exacerbate flash crashes. However, this is an overly simplistic view of how AI might be used in trading. For example, foundation models, which are trained on the same or similar data, are not currently directly suitable for developing trading strategies. Still, there are existing mechanisms in place to address the risk of crowding or market volatility following previous flash crashes and algorithmic trading incidents. Trading algorithms today operate behind a control layer with limits on volume, price and liquidity, with circuit breakers and kill switches to mitigate the impact of market volatility. Currently, Markets in Financial Instruments Directive (MiFID) II Article 17[4] and Regulatory Technical Standards 6[5] set out rules for algorithmic trading, and Article 48[6] sets out rules for circuit breakers applicable to trading venues. These mechanisms all apply regardless of the technology being used to facilitate trading.

3. EU AI Act and financial legislation requirements

3.1 Financial sector guidance

As set out throughout this response, the financial services sector is already subject to numerous regulations, including those related to the use of technology, and firms have in place risk management frameworks to ensure compliance with these rules. Firms will continue to adapt these frameworks to meet EU AI Act requirements relating to the appropriate category of AI system, including compliance

[4] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32014L0065 (most recent version dated 28/03/2024)

[5] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32017R0589

[6] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32014L0065 (most recent version dated 28/03/2024)

with rules on high-risk AI systems. For this reason, financial service sector-specific guidance is not required at this stage. Introducing further guidance for financial institutions runs the risk of increasing regulatory complexity in an already complex regulatory landscape with little benefit from a risk management perspective. Should further sector-specific guidance be introduced, it should be principles based, technology neutral and developed with substantive input from financial regulators to ensure existing governance and control frameworks are taken into account.

3.2. EU AI Act

Without further specific technical guidance at this stage, it is difficult to comment on potential gaps in the rules set out by the EU AI Act, although, as acknowledged by the Act, there are multiple areas that require further clarification through technical standards or delegated acts. One area of the Act that could benefit from further clarification is the definition of AI. The definition of AI in the Act – which is based on the Organisation for Economic Co-operation and Development (OECD) definition of AI,[7] although not exactly the same – is 'a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments'. While the concept of autonomy should be included in the definition, firms would benefit from clarification on the precise meaning of autonomy in this context.

Recital 12 states that autonomy means 'some degree of independence of actions from human involvement and of capabilities to operate without human intervention'. There is no mention, however, of whether a system can learn or act, or both, which is key to determining which AI systems are in scope; the OECD definition distinguishes between both when referring to autonomy. The use of 'may' in relation to exhibiting adaptiveness suggests an AI system could be adaptive (clarified in Recital 12 as 'a system's self-learning capabilities, allowing the system to change while in use') or not, and be captured by the definition of AI system in either case. These are just two examples where based on an initial review, firms could benefit from further clarity to determine which systems are in scope of the Act. Any guidance on the definition of AI should be led by the European Commission's Directorate-General for Communications Networks, Content and Technology in cooperation with the industry.

The Act supports regulatory sandboxes to foster AI innovation in a controlled environment by enabling the development, testing and validation of AI systems while ensuring their compliance with regulations before they enter the market. Regulatory sandboxes will provide a way to address legal uncertainties and promote evidence-based learning for regulatory authorities. The Act also allows testing AI systems in real-world conditions outside the sandboxes to accelerate the development and

[7] The OECD definition of AI is '*An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems* vary *in their levels of autonomy and adaptiveness after deployment*.'

market entry of high-risk AI systems. The challenge here will be to balance innovation with the responsible use of AI in real-world settings.

Conclusion

The efforts to better understand the use of AI in finance is a positive step, as the European Commission considers the regulatory treatment of the use of AI by financial services firms. Financial institutions have been using AI for several years in some cases and are familiar with the technology and managing the associated risks, in line with existing financial services regulation. For this reason, sector-specific guidance for financial services is not needed at this stage. As the Commission considers potential future guidance, it should take into account existing financial services legislation and risk management practices, working with the European Supervisory Authorities to ensure final guidance is proportionate and does not add regulatory complexity for firms.

A collaborative approach among financial institutions, third-party AI providers and regulatory bodies is essential for navigating the complexities of generative AI implementation. Shared responsibility ensures comprehensive risk management, aligning technological advancements with operational and regulatory needs.

Regulators should support the development of global standards and their use across the financial services and regulatory landscape by explicitly recognising such standards as presumptive evidence of compliance with existing financial industry regulations. In addition, regulators should foster industry collaboration and training based on such standards.

Given the inherent scale and borderless nature of AI technologies, cohesive, consistent and rationalised regulatory frameworks are critical. A lack of international alignment can impede companies from developing and deploying products leveraging AI.