



14 October 2026

The Next Frontier of Finance: Unlocking Agentic AI

Executive summary

Agentic AI is not a future concept but a transformation that is already underway across the economy, including in financial services. According to the World Economic Forum, the global AI agents market, valued at \$5.4 billion in 2024 and projected to reach \$236 billion by 2034, is rapidly expanding into core enterprise functions, reshaping how organisations operate, make decisions and interact with customers.¹

AI agents capable of acting on behalf of users are beginning to improve financial management, enable new digital services and streamline payments, making financial services more efficient, personalised and readily available to consumers and businesses.

At the same time, the impact of agentic AI extends well beyond retail use cases. Organisations are increasingly deploying agentic systems to augment decision-making, automate complex workflows, coordinate activities across functions and operate with a degree of autonomy in operational, risk, compliance and customer-facing processes. As adoption accelerates, agentic AI has the potential to reshape how work is performed, decisions are made and organisations govern technology. This creates significant opportunities whilst introducing new challenges related to accountability, oversight, transparency and risk management.

To support the development of agentic AI in financial services, DIGITALEUROPE calls on EU policymakers to consider the following recommendations:

- ▶ Ensure that the development of agentic AI in financial services is based on clear governance, consumer protection, liability, standards and interoperability, evaluation and testing, transparency and human accountability. Existing regulatory frameworks, including those governing financial services, should be fully implemented and their interaction with Agentic AI clarified before any new rules are considered.
- ▶ Anchor agentic finance in the principle that the user remains in control: AI agents should act only within the boundaries authorised by the user. Delegated authority and transactions carried out under it should be independently identifiable and verifiable, with authentication proportionate to the risk and consistent with existing framework.

¹ World Economic Forum, *AI agents could be worth \$236 billion by 2034 – if we ensure they are the good kind*, available at <https://www.weforum.org/stories/2026/01/ai-agents-trust/>



- 
- ▶▶ Maintain a risk-based, proportionate and technology neutral regulatory approach that focuses on specific use cases and outcomes rather than regulating agentic AI as a standalone technology category.
 - ▶▶ Ensure SME participation, reduce barriers for smaller operators and avoid prescriptive requirements that only larger firms have the resources to meet.
 - ▶▶ Promote international regulatory cooperation, the development of common standard and convergence to support interoperable governance frameworks that foster innovation whilst maintaining trust, security and consumer protection.



Table of contents

Executive summary	1
Table of contents	3
Agentic AI: the future of Europe's commerce and financial services	4
Technical features of a robust agentic AI framework	4
Definitions	4
Agentic financial services & agentic commerce	5
Governance	5
Accountability and risk management	6
Transparency & consumer protection	6
Safety & liability	7
Third-party risk management	7
Human accountability	8
Data processing and sharing	8
Standards and interoperability	9
Convergence of services and systems	10
Policy recommendations	10
The EU regulatory framework	10
Agentic AI in payments	11
Agentic AI for customer onboarding	12
Agentic AI in investment advice and execution	12
International cooperation	13

Agentic AI: the future of Europe's commerce and financial services

Agentic AI represents a significant paradigm shift, requiring financial institutions to consider both how they can use AI and how AI agents will be used to access financial services. Both perspectives offer opportunities to make financial services more efficient, personalised and accessible. Policymakers should therefore seek to enable, rather than hinder, the development and deployment of agentic AI solutions across the financial services ecosystem.

Agentic AI can only deliver its full potential if it is grounded in trusted, high-quality and appropriately permissioned data and operates under clear mandates. Access to reliable and authoritative data sources is essential to ensure that AI agents operate accurately, securely and in compliance with sector-specific requirements. Equally important is the effective application and enforcement of clearly defined governance mechanisms, including existing frameworks that ensure consumer protection, liability and transparency. We also encourage the development of interoperable standards governing how AI agents interact with the broader financial ecosystem to perform tasks or orchestrate customer journeys.

These foundations are particularly relevant in financial services, where a comprehensive regulatory framework already exists and decisions and actions often depend on precise calculations, robust governance frameworks, clear data provenance and the ability to explain outcomes to customers, regulators and other stakeholders. As the ecosystem is expected to evolve towards interactions between AI agents provided by financial services firms and independent third-party agents, it will be essential to ensure that consistent standards for transparency, security and accountability apply across all providers and solutions. This will help mitigate risks, foster interoperability and build trust among consumers, financial institutions and public authorities.

Technical features of a robust agentic AI framework

Providing a clear definition of agentic AI is essential to laying the foundation for the development and adoption of this innovation in financial services. Similarly, the successful deployment of agentic AI depends on a robust framework encompassing governance, liability, consumer protection and standards that support interoperability, transparency and accountability.

Definitions

According to the OECD, agentic AI 'refers to systems composed of multiple co-ordinated AI agents that can break down tasks, collaborate and pursue complex objectives autonomously over extended periods. These systems are designed to operate in more open-ended, less predictable physical or virtual environments and to function with minimal human supervision.'² The OECD also distinguishes agentic AI from individual AI agents by describing them 'as a class of systems capable of perceiving and acting upon their environment with a degree of autonomy, employing tools to accomplish defined objectives whilst adapting to changing input and operating conditions'.³

² OECD, *The Agentic AI landscape and its conceptual foundations*. OECD Artificial Intelligence Papers. February 2026. No.56 available at https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/02/the-agentic-ai-landscape-and-its-conceptual-foundations_a9d4b451/396cf758-en.pdf

³ Ibid



For the purposes of this paper, references to agentic AI encompass both individual AI agents and systems of coordinated agents, with governance requirements reflecting their respective capabilities, autonomy and risks.

Agentic financial services & agentic commerce

Agentic AI in financial services encompasses a broad range of use cases and levels of autonomy, requiring a clear distinction between agentic commerce and agentic financial services activities, such as investment advice, credit decisions and insurance selection.

Public attention currently focuses primarily on agentic commerce deployments that support consumers with product discovery, comparison, shopping carts and checkout, including payments, whilst preserving their ability to review and authorise transactions. Fully autonomous agentic commerce transactions, in which agents independently select merchants and products and commit funds without real-time confirmation, remain a longer-term prospect in Europe. Regulatory and governance requirements should therefore be proportionate to the risks and degree of autonomy of each use case and avoid applying requirements designed for fully autonomous transactions to assisted interactions that do not materially alter consumers' risk exposure. At the same time, policymakers should begin assessing how regulatory and governance frameworks can accommodate fully autonomous transactions with appropriate safeguards.

Recent developments also point towards the use of AI agents in securities markets and “agentic investing”. As with agentic commerce, these applications are likely to involve varying levels of autonomy, depending on the agents' design and the scope of their mandates.

Beyond consumer-facing contexts, agentic systems are also being deployed in wholesale and market-infrastructure workflows. Examples include agents supporting execution benchmarking and transaction-cost analysis across the trade lifecycle, assembling research and analytical outputs with verifiable data lineage and assisting with financial-crime screening and ongoing monitoring. These applications generally operate on governed, licensed datasets within defined access entitlements, serve professional users in regulated firms and do not involve consumer funds. Their risk profile therefore differs materially from that of consumer-facing agentic commerce and requires a correspondingly proportionate governance approach.

Governance

The governance of agentic AI should be risk-based, cover the full lifecycle and be proportionate to the specific context in which it is deployed, whilst supporting innovation. Policymakers should avoid treating agentic AI as a standalone category and instead focus on the risks and outcomes associated with particular use cases within their relevant ecosystem. This approach enables organisations to build on existing governance frameworks whilst avoiding the creation of duplicative or overlapping frameworks.

Policymakers should assess whether firms have appropriate governance arrangements in place, including controls over data quality and access rights, effective human oversight, robust testing and validation procedures, ongoing monitoring of agent behaviour and performance and auditable records of agents' decisions and actions, including the provenance of underlying data. Given that agentic systems are non-deterministic and may change as underlying models, tools and data sources evolve, testing and validation should be continuous. Firms should define containment and escalation procedures in advance for situations in which critical controls fail or authorised operating boundaries are exceeded. They should also provide auditable evidence of performance based on repeatable evaluation methods and clearly defined criteria. Supervisory guidance should clarify what constitutes acceptable minimum levels of adequate evidence so





that firms are not left to infer evaluation expectations from general principles. Such criteria should be consistent with existing governance requirements.

Governance obligations should be allocated across the lifecycle according to each actor's role and level of control. Responsibilities for an agent's operation should be allocated across the ecosystem according to each actor's role, applicable legal obligations and effective control over the relevant function or decision point. Which actor is best placed to identify and mitigate risks and implement controls may vary between use cases and activities. This allocation should reflect each actor's ability to prevent or mitigate harm, without displacing responsibilities established by applicable law, and should be supported by proportionate transparency between actors so that each can fulfil its obligations. Proportionate transparency from upstream actors should ensure that deployers have the information needed to discharge their responsibilities.

Deployment-stage obligations should apply to developers of open-source models, frameworks and toolkits only insofar as they exercise relevant control over the agent's deployment and operation. Governance requirements should also reflect the agent's risk profile, including its delegated authority, autonomy and potential impact, rather than the technology employed.

Accountability and risk management

Consumer protection, duty of care and fraud prevention must remain at the core of the deployment of agentic AI in financial services and commerce. Safeguards proportionate to the different risks posed by AI applications should be embedded into the design of protocols and products from the outset and applied consistently across every party that accesses, processes or acts on consumer data and funds. Ensuring a high and consistent level of protection throughout the ecosystem will be essential to maintaining consumer trust and mitigating fraud risks.

Transparency & consumer protection

Transparency requirements for agentic AI in financial services should be meaningful, context-specific, workable and fully consistent with the framework established by the AI Act.⁴ Consumers should understand when they are interacting with an AI agent, what services it provides, what decisions and actions it is authorised to take on their behalf, the scope of its delegated authority and how to grant or revoke it. They should understand the extent of their liability and, at an appropriate level of detail, how relevant decisions are made.

The content and scope of transparency requirements should be calibrated to avoid information overload. Existing information and disclosures provided by companies deploying and operating the relevant models should be leveraged, to reduce duplication throughout the lifecycle. At the same time, transparency obligations should not require disclosure of sensitive fraud controls, security methods or model details, where doing so could undermine consumer protection or operational resilience.⁵

⁴ Regulation (EU) 2024/1689

⁵ Relevant AI Act requirements that should apply to agentic AI in financial services include obligations under Art 50 for providers or deployers offering an AI agent interacting with end-users and where an agent relies on an underlying general-purpose AI (GPAI) model, the requirements for GPAI model providers under Arts 53 and 55.

Safety & liability

Agent access to financial services accounts should be channelled through secure and permissioned interfaces that provide clear controls, transparency and accountability, whilst preserving existing security and authentication safeguards.

Where an agent initiates or transmits a payment order, that function should be performed by, or through, an appropriately authorised payment service provider (PSP) and be subject to the applicable payments framework. Alternatively, an agent may communicate a user instruction to a regulated PSP that performs the initiation. In all cases, policy should support secure mechanisms such as agent identity, tokenisation, consent capture, traceability, monitoring and auditable records, enabling the ecosystem to identify agent activity separately from customer actions and distinguish authorised agent activity from unknown or malicious bot activity.

In this context, liability frameworks should be based on objective and clearly identifiable factors, reflecting the respective roles of consumers, merchants, agent operators and PSPs. Responsibility should lie with the party exercising effective control over the specific function or decision point at which a failure occurred and having the ability to influence or mitigate a given outcome. This assessment should take account of AI agent's function, the degree of autonomy and ability to initiate, influence or execute financial decisions and transactions.

Financial institutions should not be held liable for AI agent actions beyond their effective control, provided they have fulfilled their own obligations. Existing refund, dispute and chargeback mechanisms offer a starting point for current agentic commerce use cases. As more advanced use cases emerge, industry should assess whether further clarification, standards or governance mechanisms are needed to apply these liability principles where an agent exceeds its delegated authority or multiple agents interact across organisational boundaries.

Third-party risk management

Where banks offer customers agentic AI developed, operated or maintained by third-party providers, existing third-party risk-management frameworks provide a strong foundation for managing the associated risks.

However, agentic AI may expose over time limitations in traditional, point-in-time oversight and assurance models that rely heavily on questionnaires, attestations and periodic reviews. These frameworks often assume stable suppliers, clearly defined boundaries and periodic oversight, whereas AI ecosystems are becoming increasingly dynamic, autonomous and interconnected. Third-party risk management will therefore need to evolve through more scalable approaches and enhanced capabilities that provide deeper insights and real-time visibility into third parties. Such approaches may include continuous monitoring of agent behaviour, control effectiveness and authorised operating boundaries, supported by predefined escalation and containment actions where material deviations are detected. As AI capabilities continue to advance, supervisory attention may increasingly shift from the existence of controls to evidence that deployed systems remain governable in practice. This includes demonstrating effective oversight of agents and the ability to intervene and contain their actions when they exceed authorised boundaries or critical controls fail.



Principles-based regulatory approaches that focus on outcomes and underlying risks, rather than technology-specific requirements, will enable both deployers and providers to modernise the mechanisms used to demonstrate effective risk management and controls.

These considerations apply where a financial services provider relies on an agent provider, directly or through its third-party service arrangements, to deliver its services. Where a customer independently uses a third-party agent outside those arrangements, the same third-party risk-management approach would not generally extend to that agent provider. In such cases, the financial services provider should not be held responsible or liable for the agent's actions beyond its effective control.

Human accountability

Human accountability should remain a central principle in the governance of agentic AI. At the same time, regulatory frameworks should recognise that effective oversight can take different forms depending on the risk profile and context of the use case.

A proportionate approach should focus on ensuring meaningful human accountability commensurate with the extent to which decisions are delegated to AI agents, whilst allowing for appropriate levels of automation.

Regulatory frameworks and supervisory guidance should recognise a spectrum of oversight models rather than a single default:

- ▶ Human-in-the-loop arrangements, in which a person reviews and approves individual actions before they take effect.
- ▶ Human-on-the-loop arrangements, in which a person monitors agent activity in real time and can intervene or halt it.
- ▶ Human-over-the-loop arrangements, in which accountability is exercised through design, mandate-setting, thresholds, escalation triggers and periodic assurance.


Distinguishing between these models prevents human-in-the-loop review from becoming the assumed baseline for every use case.

Persons and entities accountable for an agent should be able to understand the actions it has taken and assess its performance. Meaningful accountability increasingly depends not only on understanding outcomes but also on maintaining the ability to observe, intervene and, where necessary, contain or reverse autonomous actions. Oversight requirements should therefore be supported by adequate technical capabilities – including logging, traceability and performance measurement – rather than relying solely on human sign-off, which risks becoming a nominal exercise as the number of deployed agents increases.

Whilst organisations deploying agentic AI retain responsibility for how these systems are used and monitored, AI providers also have a critical accountability role, particularly in embedding security, safety, reliability and controls by design. As agentic AI capabilities evolve, organisations and regulators should systematically reassess which activities genuinely require human intervention and which continue to rely on human involvement primarily due to historical practice rather than actual risk considerations.

Data processing and sharing

Agentic architectures raise questions about how existing data-licensing and entitlement frameworks apply when an agent accesses third-party data on a user's behalf. In financial services, a substantial proportion of the data on which agents rely is licensed under terms defining who may access it, for what purpose and





with which redistribution rights. Access granted to a human user should not be assumed to extend automatically to an agent acting on that user's behalf or to outputs generated by that agent.

Licensing practices and industry standards are evolving to support machine-readable entitlements, provenance and lineage mechanisms that allow data rights to be verified and enforced when an agent accesses data. This protects rights holders and is also a precondition for explainability and auditability: a firm cannot demonstrate how an agent reached an outcome if it cannot establish which data the agent used and on what basis it was entitled to use it.

From a policy perspective, data shared between participants in the agentic value chain should also be proportionate to the purpose for which it is needed, consistent with data minimisation principles under the General Data Protection Regulation (GDPR) and with other relevant frameworks, including sectoral data-access and open banking rules, whilst enabling the responsible deployment of agentic AI.⁶

Agentic payments may nevertheless require the processing of additional information to demonstrate customer authorisation. Relevant frameworks should ensure that the information necessary to support secure and trusted agentic transactions is available whilst preserving privacy and data-protection principles.

Particular consideration should be given to purpose limitation, transparency, retention, onward sharing, the withdrawal of permissions and the distinction between access permissions, payment authorisation and the legal basis for processing personal data under the GDPR.

Responsibilities among data controllers and processors across the value chain should be allocated according to each actor's legal role and level of control, without automatically placing sole responsibility on financial institutions.

Standards and interoperability

Robust ecosystem standards will be critical to enabling the safe and scalable deployment of agentic AI in financial services.

Agentic AI should not develop as a series of closed and proprietary ecosystems. A key policy objective for Europe should be to support open, interoperable and secure interfaces that enable firms to connect AI agents to trusted data sources, tools and services without locking customers into a particular model, cloud provider or vendor ecosystem.

Open standards can promote competition, innovation and portability whilst reducing integration costs and fragmentation, particularly for smaller players. They can also promote transparency and appropriate controls across the value chain, enabling all participants – including consumers, merchants, financial institutions and agent operators – to understand and manage how agents interact with their systems and data. Such standards are also relevant to digital operational resilience, as addressed under the Digital Operational Resilience Act (DORA), and to digital sovereignty, as dependence on a limited number of non-EU foundation models or agent infrastructures can increase concentration and third-country data transfer risks.⁷

Emerging open protocol frameworks can help address challenges related to integration standardisation, governance boundaries and interoperability across providers, whilst supporting secure and controlled access to financial systems and data. However, open standards should extend beyond connectivity to

⁶ Regulation (EU) 2016/679

⁷ Regulation (EU) 2022/2554



encompass observability. Common formats for agent traces, evaluation results and audit records would allow firms to compare, monitor and evidence agent behaviour across providers whilst preventing evaluation and audit data from becoming a new source of vendor lock-in. Financial institutions and data providers are already deploying model-agnostic protocol interfaces of this kind to deliver content into multiple agent environments without tying customers to a single model or cloud provider, which demonstrates that interoperability is achievable through industry practice rather than regulatory prescription.

Policymakers should avoid prescribing specific technical architectures too early and should instead support industry-led standards for agent identity, secure interoperability, authentication, auditability and dispute resolution. At the time of writing, standards for agentic AI are still at an early stage of development. Appropriate frameworks are likely to emerge through iterative testing and learning as the landscape matures. Over time, this process should support industry convergence towards common, non-proprietary approaches that strengthen interoperability and lower barriers to entry.

Convergence of services and systems

The convergence of different services and systems is increasingly evident as AI providers move into advice-like activities. In this context, a same service, same risk, same regulation approach is important, particularly where AI platforms are not only providing advice-like outputs but are also capable of taking actions on behalf of users. This principle should apply when AI platforms enter the realm of regulated financial services – such as investment advice or credit assessment – but not to routine commercial activities in which agents support product discovery and checkout within existing relationships between merchants and payment service providers.

As the distinction between advice and execution becomes increasingly blurred, the potential impact on customers and the associated risks may be similar regardless of whether an AI-driven action is performed by a regulated financial institution or an external technology platform. Without consistent regulatory treatment, comparable risks could be subject to different standards, creating regulatory gaps, competitive distortions and potential risks to consumers and financial stability.

Policy recommendations

The regulatory approach to agentic AI in financial services should focus on legal certainty, supervisory dialogue and international alignment and should be based on a thorough assessment of existing frameworks to prevent regulatory overlap and inconsistencies.

The EU regulatory framework

The EU should not reopen the AI Act or introduce new Agentic-AI-specific financial services legislation. Instead, policymakers should build on existing European frameworks and sector-specific financial services legislation, payment security requirements, consumer protection rules and existing contractual and liability frameworks, whilst clarifying how they apply to agentic use cases and addressing targeted gaps where necessary. These frameworks include the AI Act, DORA, the GDPR, the electronic Identification, Authentication and trust Services (eIDAS) Regulation, the Payment Service Directive 2 (PSD2) and the forthcoming Payment Service Directive 3 (PSD3) and Payment Service Regulation (PSR).⁸ In particular, policymakers should clarify how the AI Act's existing documentation and record-keeping obligations apply

⁸ Regulation (EU) 2022/2554, Regulation (EU) 2016/679, Regulation (EU) 910/2014, Directive (EU) 2015/2366, COM(2023) 366 final and COM(2023) 367 final, respectively.



to agentic systems, including which actor is responsible for documenting each stage of a multi-agent chain. This would allow firms to build on the technical documentation they already maintain rather than create parallel processes.

To support innovation and responsible adoption, regulators should promote structured dialogue between industry and supervisors, ensure consistent and harmonised application of legislation across Member States and sectors, use regulatory sandboxes and testing environments where appropriate and develop guidance that is technology-neutral, outcomes-based and internationally interoperable.

This approach would enable the ecosystem to identify where existing regulatory frameworks, supervisory expectations and market infrastructures remain appropriate and where targeted clarifications or guidelines might be necessary to accommodate emerging agentic finance use cases.

Agentic AI in payments

Agentic AI in payments enables software agents to perform payment-related actions across discovery, initiation, execution and orchestration, based on instructions and constraints defined by the payment service user (PSU). This raises questions about the application of strong customer authentication (SCA) requirements under PSD2 and the forthcoming PSR, which were largely designed around user-present authorisation.

Existing payment frameworks may accommodate automated execution in agentic payment flows by building on established approaches. Agentic payments should be permitted where the customer and agent can be identified, the consumer's authorisation can be reliably established and the transactions can be confirmed as falling within its scope, without mandating a specific verification mechanism. Delegated authority should be explicit, limited, revocable and auditable. Agents should not access payment accounts through credential sharing or impersonation. PSPs should retain the ability to assess transactions, suspend access and preserve sufficient evidence.

The framework should distinguish between agents provided by account servicing payment service providers (ASPSPs) and third-party agents. Where the activities of a third-party agent constitute payment initiation service, the agent should operate as, or in conjunction with, a payment initiation service provider (PISP). Third-party agents should identify transactions as agent-initiated and provide evidence that each transaction falls within the scope authorised by the user and supply ASPSPs and, where relevant, acquirers with the information necessary to apply the authentication, transaction monitoring and fraud-prevention measures for which they are responsible. The framework should also clarify which participant is responsible for verifying that the agent acts within the scope of the user's delegation, without altering the allocation of consent responsibilities under the applicable payment services legislation.

The Regulatory Technical Standards on SCA under Art. 89 of the PSR could clarify, following a risk-based approach, how existing requirements apply to agentic payments.

Authentication requirements should be capable of adapting to new technological solutions, whilst remaining proportionate to the transaction risk and consistent with existing risk-based frameworks. This includes Transaction Risk Analysis (TRA) exemptions under PSD2, transaction monitoring mechanism and the risk assessment requirements provided for in Art. 85 of the PSR.

As agentic payments mature, regulatory sandboxes could provide a controlled environment to explore how SCA requirements can be met, including whether SCA performed when an upfront mandate is established could cover subsequent transactions within its parameters and when TRA exemptions may be appropriate. Any such mandate should be clear, specific, revocable and subject to appropriate safeguards and defined





boundaries, such as duration, transaction amounts and types, beneficiaries or merchant categories. These experiments should also assess whether and under what conditions the upfront mandate could constitute the user's consent to subsequent transactions within those boundaries, whilst distinguishing consent from authentication requirements. Liability should be governed by applicable law and reflect each party's effective control over the relevant function or decision point. Legal certainty should also be provided regarding refunds, complaints and dispute resolution.

Broader liability questions, including agent malfunction, erroneous recommendations or actions outside the user's mandate, should be addressed within the relevant underlying relationships, including those between agent providers and users or between consumers and merchants, under applicable consumer protection, AI, product liability or contract law frameworks. PSPs should not be held responsible for failures of an independent agent or a merchant-controlled environment where they lack effective control over the relevant function or decision point.

Agentic AI for customer onboarding

The use of AI agents should not alter existing anti-money laundering (AML) and Know your Customer (KYC) obligations.⁹ Where an AI agent participates in account opening or customer onboarding, the natural person or legal entity behind the delegation or mandate must remain identifiable, whilst the regulated entity must remain responsible for customer due diligence, identity verification and AML checks. Agent-assisted onboarding should build on trusted digital identity mechanisms, including eIDAS and European Digital Identity Wallets. An agent should not itself be the legal subject of identity proofing but should present verifiable credentials linking it to the customer or authorised representative on whose behalf it acts. Consideration may also be given to developing "Know Your Agent" principles that complement existing KYC requirements by providing transparency regarding the identity, provenance and authority of agents acting on behalf of customers.

Agentic AI in investment advice and execution

Investment advice and execution are both potential applications of agentic AI and the risks associated with agent-led execution warrant specific consideration. Where an agent selects and purchases an investment product on a user's behalf, Market in Financial Instruments Directive II (MiFID II) suitability and appropriateness obligations should continue to apply where relevant.¹⁰ Policymakers should clarify how the client is identified when an agent transacts on their behalf, whether an upfront mandate can satisfy the suitability assessment and how compliance with appropriateness requirements can be demonstrated when the client is not present at the point of execution.

As the distinction between regulated investment advice and comparable technology-enabled services may become increasingly blurred, regulatory classification should reflect the substance of the activity, including whether the system exercises discretion, provides a personal recommendation, initiates transactions or influences customer decision-making in a manner comparable to regulated advice.

⁹ Regulation (EU) 2024/1624

¹⁰ Directive (EU) 2014/65



The principle of “same service, same risk, same regulation” should apply to execution as well as advice-like outputs. Comparable clarification is needed for agent-mediated insurance distribution under the Insurance Distribution Directive (IDD), including the application of the demands-and-needs test.¹¹

International cooperation

Given the global nature of AI development and deployment, policymakers should promote international regulatory cooperation and convergence on the governance of agentic AI. Greater alignment across jurisdictions can help reduce regulatory fragmentation, provide legal certainty for innovators and support the development of interoperable governance frameworks.

Encouragingly, regulators in other jurisdictions have already begun exploring these issues. For example, in March 2025, the UK's Digital Regulation Cooperation Forum (DRCF) published a discussion paper on the future of agentic AI, setting out shared definitions, key regulatory considerations and areas for further collaboration.¹²

Work is also advancing at the international level. The International Organisation of Securities Commissions (IOSCO) has published supervisory tools for the use of AI in capital markets.¹³ Furthermore, the Financial Stability Board (FSB) has developed a catalogue of sound practices for responsible AI adoption that addresses issues, vulnerabilities and dependencies associated with agentic AI across the financial system.¹⁴ The Organisation for Economic Co-operation and Development (OECD) also provides a natural forum for promoting interoperability among national approaches.¹⁵

As these bodies shape supervisory expectations well before they are reflected in national rules, the EU should engage actively in these forums to promote international consistency and prevent regulatory divergence, particularly where overarching principles already converge.

¹¹ Directive (EU) 2016/97

¹² Digital Regulation Cooperation Forum, *The Future of Agentic AI. Foresight paper* available at <https://www.drcf.org.uk/siteassets/drcf/pdf-files/drcf-the-future-of-agentic-ai-foresight-paper.pdf?v=416157&ct=1781261193692>

¹³ IOSCO, *Supervisory Toolkit for AI Use in Capital Markets*, available at <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD823.pdf>

¹⁴ FSB, *Sound Practices for Responsible Adoption of Artificial Intelligence (AI). Consultation report*, available at <https://www.fsb.org/uploads/P100626.pdf>

¹⁵ OECD, *Artificial Intelligence Papers. Agentic AI in organisations: Early insights from practitioner interviews*. No.65, available at https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/09/agentic-ai-in-organisations_1f17d235/1257a26f-en.pdf



FOR MORE INFORMATION, PLEASE CONTACT:

Federico Di Benedetto

Senior Manager for Digital Transformation of Financial Services

federico.dibenedetto@digitaleurope.org / +32 493 12 34 88

Vincenzo Renda

Director for Digital Transformation Policy

vincenzo.renda@digitaleurope.org / +32 490 11 42 15

About DIGITALEUROPE

DIGITALEUROPE is the leading trade association representing digitally transforming industries in Europe. We stand for a regulatory environment that enables European businesses and citizens to prosper from digital technologies. We wish Europe to grow, attract and sustain the world's best digital talents and technology companies. Together with our members, we shape the industry policy positions on all relevant legislative matters and contribute to the development and implementation of relevant EU policies. Our membership represents over 45,000 businesses who operate and invest in Europe. It includes corporations which are global leaders in their field of activity, as well as national trade associations from across Europe.

