© RELX (UK) Limited trading as LexisNexis.
Not for external distribution or resale

Feature

KEY POINTS

- Artificial Intelligence (AI) regulation in financial services is evolving from conduct risk through governance, operational risk, prudential risk and towards financial infrastructure. Agentic AI – systems operating without continual human intervention – has accelerated this transition towards prudential and infrastructure frameworks, with profound implications for standard ISDA, GMRA and GMSLA documentation.
- Representations, warranties, events of default and termination triggers in standard trading documentation must now contemplate AI system failures – whether occurring at the borrower or lender – that cascade across interconnected operational functions.
- The emerging regulatory paradigm addresses AI-dependent financial systems – requiring infrastructure-level supervision and new models of supervisory engagement.

Author Dr Michael Huertas

AI in the trade lifecycle: documentation implications in light of the five stages of AI regulation in the EU and UK

This article traces five regulatory stages for AI in financial services: conduct risk, governance, operational risk, prudential risk and financial infrastructure. Agentic AI has accelerated the transition towards prudential and infrastructure frameworks, shifting regulation from addressing AI itself to addressing AI-dependent financial systems – a fundamentally different undertaking with profound implications for trading documentation.

INTRODUCTION

AI regulation in financial services is following a trajectory comparable to that of operational risk after the Global Financial Crisis – from a conduct of business to a prudential issue towards market infrastructure. This has potentially profound implications for trading documentation. Representations, warranties, events of default and termination triggers in ISDA, GMRA and GMSLA documentation were drafted for a world in which technology was a back-office function, not a determinant of prudential soundness. That assumption is increasingly difficult to sustain.

This article traces five stages in the regulatory treatment of AI and argues that agentic AI – having achieved early success in payments – is now advancing into trade execution, collateral management and regulatory reporting. At each stage, the implications for trading documentation deepen.

Before examining the five stages, it is necessary to distinguish between traditional AI (machine learning, optimisation) and generative AI (large language models). The back office has used traditional AI for over a decade; generative and agentic AI are beginning to transform those functions, though transformation remains in early stages within regulated financial institutions.

AI ACROSS THE TRADE LIFECYCLE: CURRENT STATE AND EMERGING APPLICATIONS

Across a modern regulated firm's trade lifecycle, traditional machine learning has achieved high penetration in trade matching, collateral management and margin calculation. Industry surveys from 2025-2026 suggest generative AI adoption is highest in client reporting and regulatory reporting, while agentic AI remains largely emerging – with advanced deployments appearing in exception management (managing errors) and reconciliations.

Generative AI is deployed in trade exception management, ISDA confirmation review, collateral disputes, regulatory reporting (EMIR, SFTR, MiFIR) and documentation review. Agentic AI extends further: rather than merely summarising, an AI agent can orchestrate an entire failed settlement workflow – from identifying the failure through to preparing the regulatory notification – without repeated human prompts.

Certain activities remain highly deterministic (responding to human prompts): pricing engines, VaR calculations and CCP margin models rely on established quantitative models where regulators expect transparency. Generative AI is not replacing these models.

The documentation implications flow directly from this operational reality. Trade affirmation, settlement, collateral, liquidity and regulatory reporting increasingly form a continuous AI-dependent chain. Its failure becomes operational resilience risk, concentration risk and potentially prudential risk. Representations, warranties, events of default and termination triggers must contemplate scenarios in which AI system failures – at either counterparty – cascade across interconnected functions.

There is a deeper insight: generative AI is not primarily automating financial decision-making; it is automating the orchestration of financial operations. Traditional models answered: "What is today's variation margin?". Agentic AI increasingly answers: "How should the institution organise the end-to-end workflow to calculate, verify, communicate, dispute, settle and report today's variation margin?".

THE FIVE STAGES OF AI REGULATION IN FINANCIAL SERVICES

Stage 1: AI as conduct risk (2018-2022)

The earliest regulatory engagement with AI was anchored in conduct risk – shaped by

© RELX (UK) Limited trading as LexisNexis. Not for external distribution or resale

GDPR's provisions on automated decision-making, MiFID II suitability requirements and the emerging AI Act in the EU, and the FCA's Consumer Duty in the UK. The supervisory question was: was the customer treated fairly?

For trading practitioners, the conduct stage had negligible transactional impact. Standard trading documentation was drafted by reference to financial condition and contractual performance – not technological tools. That would change.

Stage 2: AI as governance (2023-2025)

The second stage shifted focus to accountability. The EU AI Act's governance requirements, DORA's ICT governance framework and ECB model risk expectations all contributed. In the UK, PRA SS1/23 on model risk management reinforced the same trajectory. The supervisory question became: who is responsible?

The governance stage generated tangible drafting consequences. Credit Support Annexes increasingly needed to contemplate counterparty AI governance arrangements. Schedule provisions may include Additional Termination Events linked to material regulatory findings on model risk. If a counterparty's AI governance framework were found materially deficient, could that constitute a condition giving rise to early termination? The standard ISDA Section 5 framework does not expressly capture regulatory findings on AI governance, yet the commercial consequences could be material.

Stage 3: AI as operational risk (2024-2026)

The third stage represents a shift from governance to resilience: can AI fail, create outages or compromise business continuity? DORA provides the most comprehensive legislative expression, with its framework for critical ICT third-party providers. The PRA's operational resilience framework and ECB cyber supervision sit within the same paradigm. Just as operational risk moved from a governance afterthought to a capital-consuming risk category, AI risk is now making the same transition.

The documentation consequences are material. DORA-related representations and Additional Termination Events are appearing in ISDA Schedules. Counterparty due diligence requires assessment of AI provider dependencies. A critical AI system outage preventing a counterparty –whether borrower or lender – from meeting margin obligations could trigger close-out cascades with systemic implications.

Stage 4: AI as prudential risk (2025-2027)

The fourth stage is where regulation is unmistakably heading. AI becomes a source of prudential risk – creating capital risk, liquidity risk and governance failures that engage capital adequacy, stress testing and recovery planning.

Agentic systems introduce autonomous decision-making, recursive optimisation and feedback loops. When deployed in credit decisioning, treasury management or collateral optimisation, they create risks qualitatively different from traditional algorithmic tools. A credit-scoring algorithm producing biased results is a conduct problem. An agentic AI system autonomously reallocating collateral in ways that amplify procyclicality is a prudential problem.

The supervisory question becomes: can AI threaten the safety and soundness of the institution? That distinction has profound consequences for capital planning, ICAAP, ILAAP and recovery plans.

For trading relationships, the prudential stage transforms counterparty credit assessment. Where a counterparty – whether bank borrower or lender – deploys agentic AI in treasury management or collateral optimisation, credit assessment must evaluate whether those systems could autonomously generate prudential risks. Thresholds and minimum transfer amounts in Credit Support Annexes take on a different character when the counterparty's capacity to meet obligations is influenced by autonomous AI decisions. The ISDA's standard close-out architecture was not designed with this scenario in mind.

Stage 5: AI as financial infrastructure (2026 onwards)

The fifth stage sees agentic AI increasingly situated inside financial infrastructure itself – market infrastructures, payment systems, clearing, settlement, the emerging Digital Euro ecosystem and MiCAR-regulated cryptoasset services.

AI is no longer simply used by firms – it becomes infrastructure upon which the financial system operates. The regulatory model shifts from one in which institutions use AI to one in which institutions depend upon AI infrastructure, which itself becomes regulated.

The transition to T+1 settlement illustrates this. Compressing the settlement cycle requires AI-driven matching and exception handling operating beyond human capacity. Where those AI systems sit inside CSDs or CCPs, they are not tools used by individual firms – they are infrastructure upon which the market depends. An AI failure in a CSD's matching engine would be a market event, not merely an operational incident.

The embedding of AI within financial infrastructure creates risks that institution-level governance cannot contain. The solution requires infrastructure-level standards and supervisory oversight operating at system level.

AGENTIC AI AND THE TRANSFORMATION OF LEGAL ANALYSIS

The five-stage evolution is accelerated by agentic AI. Ordinary AI supports human decisions. Agentic AI initiates, co-ordinates, executes, monitors and learns without continual human intervention. The legal consequences cut across virtually every area of financial services regulation.

Agentic AI is deployed across:

- credit origination;
- treasury management;
- collateral optimisation;
- market surveillance;
- AML screening;
- compliance monitoring; and
- resolution planning.

© RELX (UK) Limited trading as LexisNexis. Not for external distribution or resale

In each domain, the shift from AI supporting a human to AI acting autonomously changes the regulatory analysis.

Three scenarios illustrate the point. First, an agentic AI system managing a bank's collateral pool across multiple CCPs optimises allocation in normal conditions but, under stress, responds to margin calls by liquidating assets and competing with other banks' agentic systems for the same high-quality liquid assets – potentially amplifying procyclical price movements. No human treasury officer authorises these individual decisions.

Second, an agentic AML system that monitors transactions, investigates alerts and files suspicious activity reports without human review raises questions about compliance validity and MLRO liability. If the system fails to identify a recently designated sanctioned entity, the resulting breach is a regulatory and potentially criminal compliance failure.

Third, an agentic system monitoring intraday liquidity that autonomously draws on credit facilities during market stress – triggering cross-default provisions – could accelerate the very funding crisis it was designed to prevent, compressing the window for management intervention to irrelevance.

The implications extend across multiple frameworks. Agentic AI changes ICAAP, ILAAP, recovery plans, risk appetite statements, outsourcing risk, concentration risk and business continuity planning.

Resolution planning deserves attention. Under BRRD and the Bank of England's framework, where critical functions depend upon agentic AI, resolution authorities must satisfy themselves that AI systems will continue operating during resolution. That requires analysis of whether AI provider contracts contain resolution-resilient terms and whether AI systems can function in a post-resolution entity.

Senior manager accountability raises challenging questions. Under SM&CR, when an agentic AI autonomously initiates collateral calls or suppresses an AML alert, which senior manager bears personal regulatory responsibility? If decisions are made at machine speed, the factual basis for a "reasonable steps" defence may be difficult to establish. The question neither regime has squarely addressed is whether a senior manager can be held accountable for autonomous AI decisions – and what governance architecture makes that accountability meaningful.

THE INFRASTRUCTURE DIMENSION: DIGITAL PRUDENTIAL SUPERVISION

AI is becoming the mechanism through which compliance is achieved, monitored and enforced.

This has direct supervisory consequences. Current trajectories point towards agentic regulatory reporting in real time, a "Supervisory Digital Twin" framework comprising an AI-powered, real time simulation of the whole financial system enabling continuous monitoring, and machine-readable regulation. The future supervisory relationship may be one in which a supervisor's AI interacts with the institution's AI.

That raises questions current frameworks cannot answer. Who is accountable when the supervisor's AI misinterprets the institution's AI-generated data? How should capital requirements be calibrated for institutions whose risk profiles change in real time?

The token economy (the method of pricing for AI usage) may amplify these dynamics. As AI consumption shifts to per-token pricing, institutions depending upon agentic AI may face cost volatility that is itself prudential risk. AI is becoming a utility and may need to be regulated like one.

THE UK COMPARISON: PRUDENTIAL AI SUPERVISION WITHOUT AI LEGISLATION

The EU has chosen the legislative path: the AI Act, DORA, MiCAR, CRD VI, the forthcoming Payment Services Regulation and Digital Euro package collectively create a dense, interlocking framework addressing AI at multiple levels.

The UK has taken a different approach – relying upon supervisory expectations and existing FCA, PRA and Bank of England powers rather than AI-specific legislation. The FCA's June 2025 speech 'Rethinking regulation in the age of AI' explicitly rejects a prescriptive model in favour of outcomes-focused regulation, emphasising meaningful human oversight, explicable decisions and concentration risks when multiple firms depend upon the same AI providers. That concentration concern suggests the UK is also moving towards Stage 5.

The UK is building prudential AI supervision without AI legislation – contrasting with the EU's legislation-led model. Neither approach is inherently superior. The EU offers greater legal certainty; the UK offers greater flexibility. Yet the FCA's approach reveals tension: outcomes-focused regulation works when firms and supervisors share understanding of good outcomes, but agentic AI introduces scenarios unpredictable to both.

For firms operating across both jurisdictions, the divergence creates complexity – simultaneously complying with EU prescriptive requirements and UK outcome-based expectations that differ materially in form and documentation.

A cross-border trading relationship must navigate potentially divergent requirements in representations, warranties, event of default, Additional Termination Events and credit support documentation. An AI-related regulatory action in one jurisdiction – whether affecting borrower or lender – could trigger close-out rights, margin calls or termination events across the group's trading arrangements in other jurisdictions.

PRACTITIONER TAKEAWAYS

Until the regulatory framework catches up with the infrastructure dimension, practitioners can take concrete steps.

First, boards should ensure AI risk is integrated into ICAAP, ILAAP and recovery planning, including scenario analysis for autonomous AI behaviour.

Second, firms should map AI infrastructure dependencies against DORA's critical ICT provider framework and subject critical dependencies to resilience testing and exit planning.

Third, firms operating across the EU and UK should develop dual-track compliance

© RELX (UK) Limited trading as LexisNexis. Not for external distribution or resale

Feature

Biog box

Dr Michael Huertas, LLM, MBA, is a partner and the Global Financial Services Legal Leader for PwC Legal. He is a Solicitor-Advocate (England and Wales), Solicitor (Ireland) and a German Rechtsanwalt. His professional practice focuses on emerging regulatory issues in the Banking Union and Capital Markets Union. The author would welcome dialogue on any of the issues raised herein or in relation to his research interests.

Contact: https://www.linkedin.com/in/michael-huertas-157a788

frameworks satisfying both EU prescriptive requirements and UK outcome-based expectations.

Fourth, practitioners should engage proactively with the infrastructure dimension – including the emerging Digital Euro and tokenisation frameworks – rather than waiting for regulatory clarity.

Fifth, practitioners should review standard derivatives, repo and securities lending documentation for adequacy in an AI-dependent environment – including Additional Termination Event provisions capturing AI-related supervisory interventions at either counterparty, representations requiring identification of material AI dependencies, and close-out triggers capturing AI-specific regulatory actions. ISDA, ICMA and ISLA may wish to consider whether standard form documentation requires updating.

CONCLUSION

The legal treatment of AI in financial services is following a trajectory comparable to operational risk after the Global Financial Crisis – from conduct risk to governance, operational risk, prudential risk and infrastructure. AI is becoming not a tool that firms use but infrastructure upon which the financial system depends.

Agentic AI accelerates this transition. Systems operating without continual human intervention create risks engaging not merely conduct and governance frameworks but prudential standards, recovery planning and the stability of market infrastructure itself.

Conduct rules risk being viewed as no longer enough. Similarly, governance frameworks while necessary may increasingly be seen to be insufficient. What is emerging in both the EU and UK is a regulatory architecture addressing AI not as a compliance problem within institutions but as an infrastructure condition of the financial system that requires a different regulatory and supervisory approach plus a level of intrusion that DORA does not deliver. Firms that get ahead of this will benefit from the advantages that responsible agentic AI offers; those that do not will be rule takers. ■

Further Reading:

- AI's current and emerging impact on restructuring and insolvency (2026) 8 JIBFL 535.
- Increasing use of AI in financial services: managing the regulatory and litigation risks (2026) 6 JIBFL 405.
- Lexis+® UK: Cyber security for financial services firms – regulatory and industry initiatives – timeline.



© RELX (UK) Limited trading as LexisNexis. Not for external distribution or resale