

# AI in Trading

## A practitioners' view of the current landscape

**Spotlight Review**

**February 2026**



## About us

### Financial Markets Standards Board

Financial Markets Standards Board Limited (FMSB) is a private sector, market-led organisation created in light of the recommendations in the Fair and Effective Markets Review (FEMR) Final Report in 2015. One of the central recommendations of FEMR was that participants in the wholesale markets should take more responsibility for raising standards of behaviour and improving the quality, clarity and market-wide understanding of trading practices. Producing guidelines, practical case studies and other materials that promote the delivery of transparent, fair and effective trading practices will help increase trust in wholesale markets. FMSB brings together people at senior levels from a broad cross-section of global and domestic market participants and end-users. In Committees and Working Groups, industry experts debate issues and develop FMSB Standards and Statements of Good Practice and undertake Spotlight Reviews - like this one - that are made available to the global community of financial market participants and regulatory authorities.

### Spotlight Reviews

Spotlight Reviews encompass a broad range of publications used by FMSB to illuminate important emerging issues in financial markets. Drawing on the insight of Members and industry experts, they provide a way for FMSB to surface challenges market participants face and may inform topics for future work. Spotlight Reviews will often include references to existing law, regulation and business practices. However, they are not intended to set or define any new precedents or standards of business practice applicable to market participants.

**Find out more about the Financial Markets Standards Board at [fmsb.com](http://fmsb.com).**

# Contents

<b>Executive summary</b>	<b>4</b>
<b>Introduction</b>	<b>5</b>
<b>Defining AI in markets</b>	<b>6</b>
– Definition	
– Applying AI techniques to trading systems	
<b>Applications and risks</b>	<b>8</b>
– Use cases	
– AI risks in trading systems	
– Accountability	
<b>Assessing AI in trading systems</b>	<b>15</b>
– Contextual considerations	
– Case studies	
<b>Conclusion</b>	<b>21</b>
<b>Annex 1 - AI definitions</b>	<b>22</b>
<b>Annex 2 – Bibliography</b>	<b>24</b>

## Executive summary

While the financial industry has long been a pioneer in the use of statistical and machine learning, the application of more advanced artificial intelligence (AI) in financial markets remains relatively nascent. Advances in computational power and the increasing accessibility of AI are nevertheless driving growing adoption. Stakeholders face the challenge of balancing AI's innovation promise against the potential risks of introducing AI into their existing systems. This Spotlight Review examines use cases of AI, their risk profiles and the suitability of existing control frameworks, seeking to advance the AI debate in financial markets. While AI is increasingly deployed in non-market-facing applications, this paper focuses specifically on the use of AI in trading. The main observations, drawn from discussions with practitioners, are:

- 1. Autonomy:** Market-facing AI is typically integrated into automated, scalable and analytics-driven trading systems. These trading systems do not operate with full autonomy and remain subject to human supervision and intervention.
- 2. Use cases:** The most common use of AI is within smaller modules of larger systems, for example, analytics that assess liquidity conditions, make venue recommendations, inform pricing predictions or produce trading metrics.
- 3. Model risk:** AI enables use cases that solve increasingly complex objective functions and tasks. The risk of such AI use cases arises from the scope and complexity of a task rather than the AI technique itself.
- 4. Monitoring outputs:** As AI models become more complex and less intuitive, attempts to interpret and explain their internal decision-making processes will face practical limits. Greater focus should instead be placed on monitoring model outputs and on independent controls proportionate to model output risk.
- 5. Control frameworks:** Many risks are already well-managed by existing control frameworks such as model risk management (MRM) or real-time algorithmic trading controls. The scale or novelty of an AI use case could, however, potentially outpace specific guardrails, which need to be kept up to date.
- 6. Human accountability:** The governance of market-facing AI requires clear accountability of human coders, traders and managers for the actions and decisions of a machine, consistent with the accountability frameworks that apply to manual and electronic trading processes.
- 7. Future of AI:** Trading systems could, over a longer-term horizon, make use of advanced techniques such as generative AI or artificial general intelligence and could reach a high degree of autonomy without human supervision. This scenario, along with its more systemic risk implications or autonomous AI-to-AI interaction, is not yet a reality in financial markets.

# 1. Introduction

Over recent decades, a significant structural shift in trading has taken place from a more manual trading model towards a stronger utilisation of technology and data. AI has the potential to transform trading and markets even further, replicating human decision-making and applying it at vast scale. The efficiency gains and revenue potential could be significant. At the same time, this transition creates new challenges and potential risks for markets.

In its infancy, technology in financial markets often focused on operational inefficiencies, shifting paper processes into an electronic world or improving the productivity of the back office. Increasingly, technology was then applied directly to trading, automating parts of the execution process or integrating data and analytics into trading systems. Today, many firms rely on automated processes for most of their trading decisions and are enhancing these processes with AI. This Spotlight Review examines AI models or systems whose outputs directly or indirectly influence pricing, execution, or client-facing trading decisions (referred to here as 'market-facing AI' or 'AI in trading') and its implications for trading risk, and the associated control environment.

Some AI model risks are already well-known and well-managed while others may be new. It is important to understand whether AI alters the risk profiles of trading activities to ensure that controls are fit-for-purpose and to help calibrate appropriate responses to these risks. The paper focuses on the tangible risks that arise from current AI model use cases, rather than predicting the risks of future scenarios which may or may not play out.

This Spotlight Review adopts a practitioner perspective of AI model uses in wholesale markets. At the current time, the landscape is relatively nascent with varying levels of adoption of AI and varying levels of familiarity with AI use cases. A path defining how best to manage AI model risk still needs to be forged and industry best practices still need to emerge.

The purpose of this publication is to provide a practitioner perspective on current market-facing AI use cases, their potential risks and their context within existing control environments, focusing on how AI models are used in practice and governed within existing control and supervision arrangements. Its observations have been drawn from discussions with practitioners from across the FMSB Membership.

The paper seeks to make the discussion less sensitive to theoretical concerns and instead sharpen its focus on tangible challenges. It is intended to complement existing work by academics and industry participants, alongside the work of policymakers, whose mandate includes anticipating and preparing for future risks.

As AI and its applications continue to evolve rapidly, follow-on work will include ongoing engagement with industry stakeholders and policymakers to monitor emerging implications for model risk and automated trading controls and to support timely and effective risk management responses.

The publication is divided into three sections:

- **Defining AI in markets**
  - Definition and applying it to AI techniques in trading systems
- **Applications and risks**
  - Use cases, AI risks in trading systems and accountability
- **Assessing AI uses in trading systems**
  - Contextual considerations and case studies

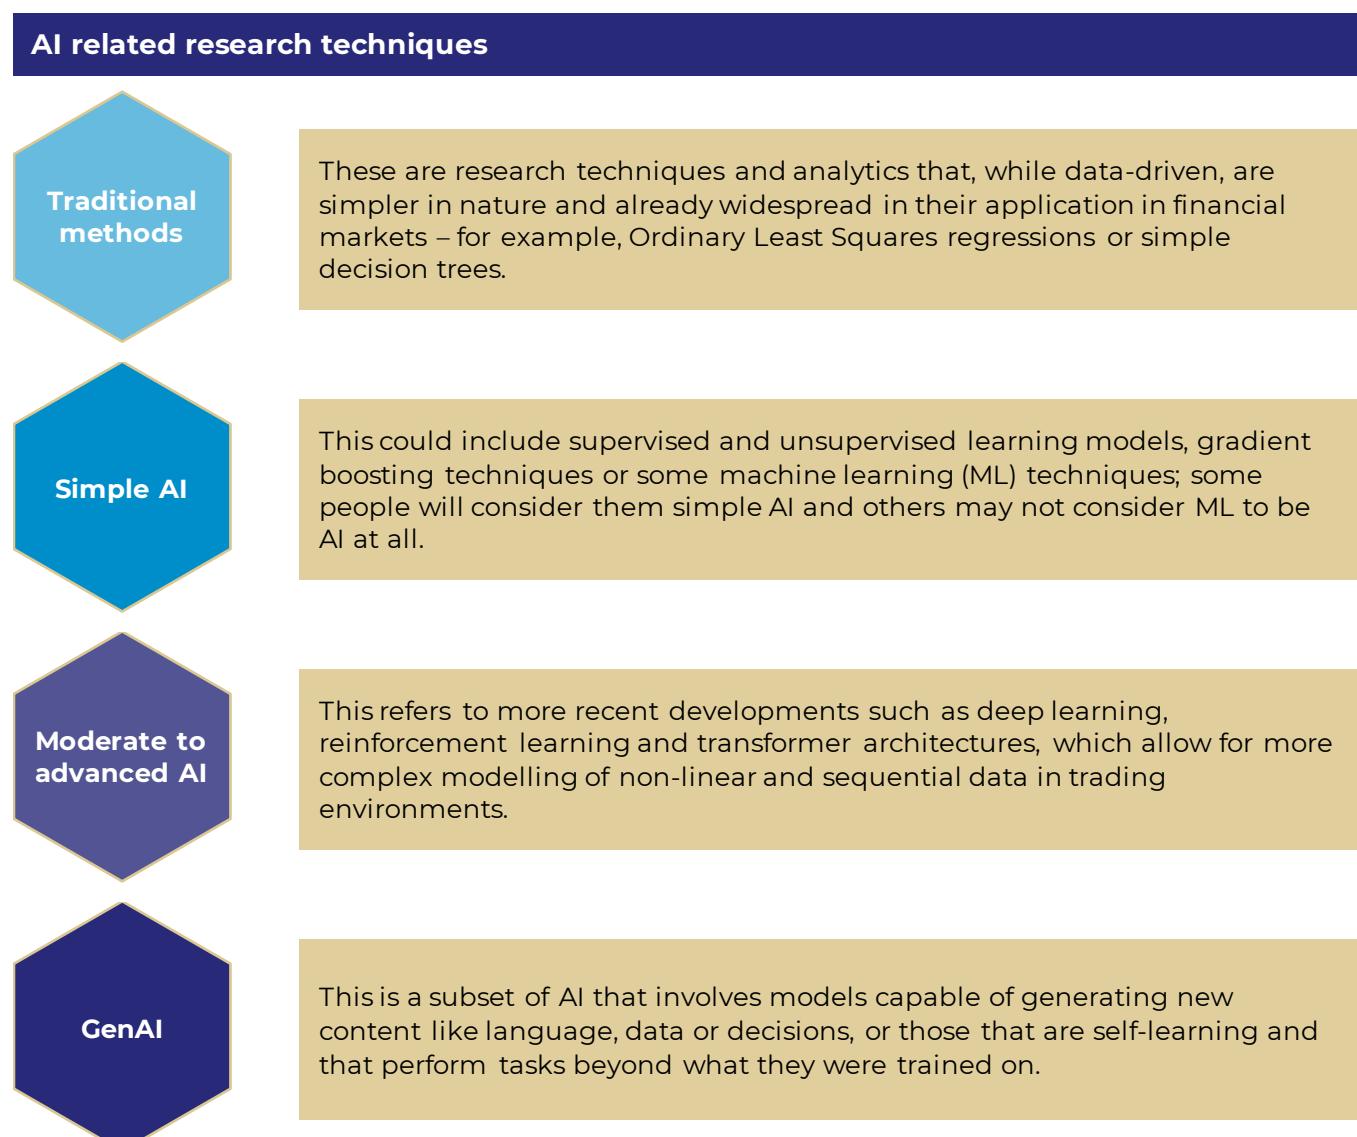
## 2. Defining AI in markets

### Definition

Defining AI in the context of wholesale markets is inherently challenging. The technology is evolving rapidly, and its applications in market environments continue to develop at pace. As a result, any definition risks becoming outdated or failing to capture relevant developments. At the same time, definitions that are framed too broadly may capture long-established quantitative techniques, creating ambiguity and potentially a disproportionate perception of risk. Annex 1 summarises existing official, non-context specific definitions, which illustrate the range of approaches currently taken.

There is a risk in drawing rigid distinctions between AI and non-AI techniques, or in developing separate AI-specific control frameworks for financial markets. Such distinctions may prove artificial or be quickly overtaken by technological change. Therefore, rather than attempting to establish a fixed boundary, this Spotlight Review focuses on a spectrum of simple to advanced AI techniques, as illustrated in Figure 1.

**Figure 1: AI research techniques**



The analysis concentrates on use cases that raise considerations distinct from those associated with well-established quantitative methods. Traditional techniques, including, for example, ordinary least squares (OLS) regressions and standard decision-tree approaches, are already embedded within existing governance and control frameworks and are therefore not considered in scope for the purposes of this paper. At the other end of the spectrum, generative AI techniques are not currently understood to be deployed in direct trading applications. Accordingly, while such techniques may warrant future consideration as technology evolves, they are not a central focus of this analysis.

## Applying AI techniques to trading systems

In trading, AI models are typically embedded within electronic trading systems that exhibit a high degree of automation, scale and reliance on analytics, enabling execution at scale. This differs from the use of AI in standalone applications, as AI models in trading environments operate within established trading systems that are subject to policies, governance arrangements and control frameworks. As a result, the risk profile of an AI model is shaped by its specific use case and the context in which it is deployed within a trading system. Risk may therefore vary significantly depending on the nature and complexity of the task being performed. While more complex tasks may necessitate the use of more sophisticated techniques, complex techniques may also be applied in lower-risk contexts.

Across current trading environments, AI models are neither deployed as standalone applications nor operated with full autonomy. Instead, they are embedded within wider electronic trading systems' procedural algorithms and are subject to independent human control and supervision layers that do not rely on AI. These control frameworks may be manual or systems-based; however, oversight and intervention ultimately rest with human traders. As a result, current market-facing AI use cases do not operate on a fully autonomous basis. Consistent with *Commission Delegated Regulation (EU) 2017/589 (RTS 6)* and other algorithmic trading rules, AI deployed in trading systems is subject to established human supervision and governance arrangements, and a transition of these supervisory functions to AI systems is not anticipated in the near term. Concerns regarding AI in financial markets are often associated with assumptions of full autonomy, which do not reflect prevailing market practice.

### 3. Applications and risks

#### Use cases

Advances in computing power have driven increased interest in AI across the global economy, including developments in large language models (LLM), deep learning, reinforcement learning, and transformer architectures. This trend has extended to the financial industry, where market participants are exploring practical AI use cases to improve efficiency and generate revenue.

Financial markets have long been early adopters of technology, data and quantitative techniques, including statistical and econometric methods used in investment strategies and algorithmic trading. Trading systems have also become increasingly automated, replacing manual processes with systems-based rules to enable execution at scale while reducing certain forms of manual risk. Building on this foundation, firms are now integrating AI into trading processes and infrastructure, with Table 1 summarising areas under exploration for AI deployment, ranging from established applications to more experimental use cases.

Many early applications of AI in institutional markets have taken the form of support tools. Using a range of AI techniques, from relatively simple to more advanced, these tools shift operational tasks from human to machine-based processes, increasing speed and scale. Despite their use of AI, such tools are typically not market-facing.

A second category of use cases employs AI as an input into larger trading systems, such as execution management systems (EMS) or algorithmic trading engines. This is currently the most commonly cited application of AI in trading. Examples include price forecasts, liquidity metrics and the modelling of liquidity or hit rates. In these cases, AI enhances data inputs and estimation techniques while remaining one step removed from execution, feeding into non-AI systems that interact directly with the market and representing an incremental improvement to long-established techniques.

A third category applies AI to elements of trading system logic, such as order routing systems that use AI analytics to determine venue selection, split parent orders, or select instruments for execution. These approaches are not technically new and were, in many cases, already deployed at scale prior to recent advances in AI.

**Table 1: AI deployment in wholesale markets – sample use cases**

AI focus	Sample use cases	Potential AI benefits
<b>Support tools</b>	<ul style="list-style-type: none"> <li>AI-based research report generation or review of trade documentation</li> <li>Analysis of client communication with natural language processing (NLP) to identify client preferences</li> <li>Processing of unstructured trade requests (chat/phone) with NLP to extract key details and create trade tickets</li> </ul>	<ul style="list-style-type: none"> <li>Shift from human to machine-based reviews increases speed and scale</li> <li>Wider and more comprehensive range of inputs</li> </ul>
<b>Input modules</b>	<ul style="list-style-type: none"> <li>Liquidity metrics or further trading signals</li> <li>Price and volume forecasts for a trading algorithm</li> <li>Trend analysis and metrics displayed in an EMS</li> <li>Forecasts at what prices to make markets</li> <li>Modelling central limit order books or hit rates</li> </ul>	<ul style="list-style-type: none"> <li>New or enhanced input data</li> <li>More powerful and flexible forecasting techniques</li> </ul>
<b>System logic</b>	<ul style="list-style-type: none"> <li>Smart order router (SOR) to optimize execution across venues or split parent orders</li> <li>Risk management of transactions and hedging costs</li> <li>Trading bots with the power to act on behalf of persons or to execute more complex sequences of tasks</li> </ul>	<ul style="list-style-type: none"> <li>Potential for automating sequences of tasks that include decision-making without human involvement</li> <li>Ability to further scale electronic trading</li> </ul>

## AI risks in trading systems

Trading systems have evolved to operate at high levels of automation and scale. As market participants shift manual processes into more electronic processes, they seek to improve capacity, reduce manual risks and grow revenues. While manual risks have been reduced, the transition has introduced different risks such as model risk or technology risk which were less prevalent in manual processes. Firms have at the same time adopted new control frameworks to manage these different risk profiles.

### Model risk

The introduction of AI into trading systems has the potential to further impact the risk profile of trading activities. In some cases, AI models will be limited in their impact and may not affect trading risk at all. In other cases, AI models may heighten existing risks or introduce new risks which were not present before. The impact of AI on trading risk, and how well it is already managed, will depend on the specific use cases. It will be important to ensure that with the introduction of AI, model risk controls and algorithmic trading guardrails keep pace with the scale and complexity of the trading systems.

**Figure 2: AI use cases by complexity and scope**

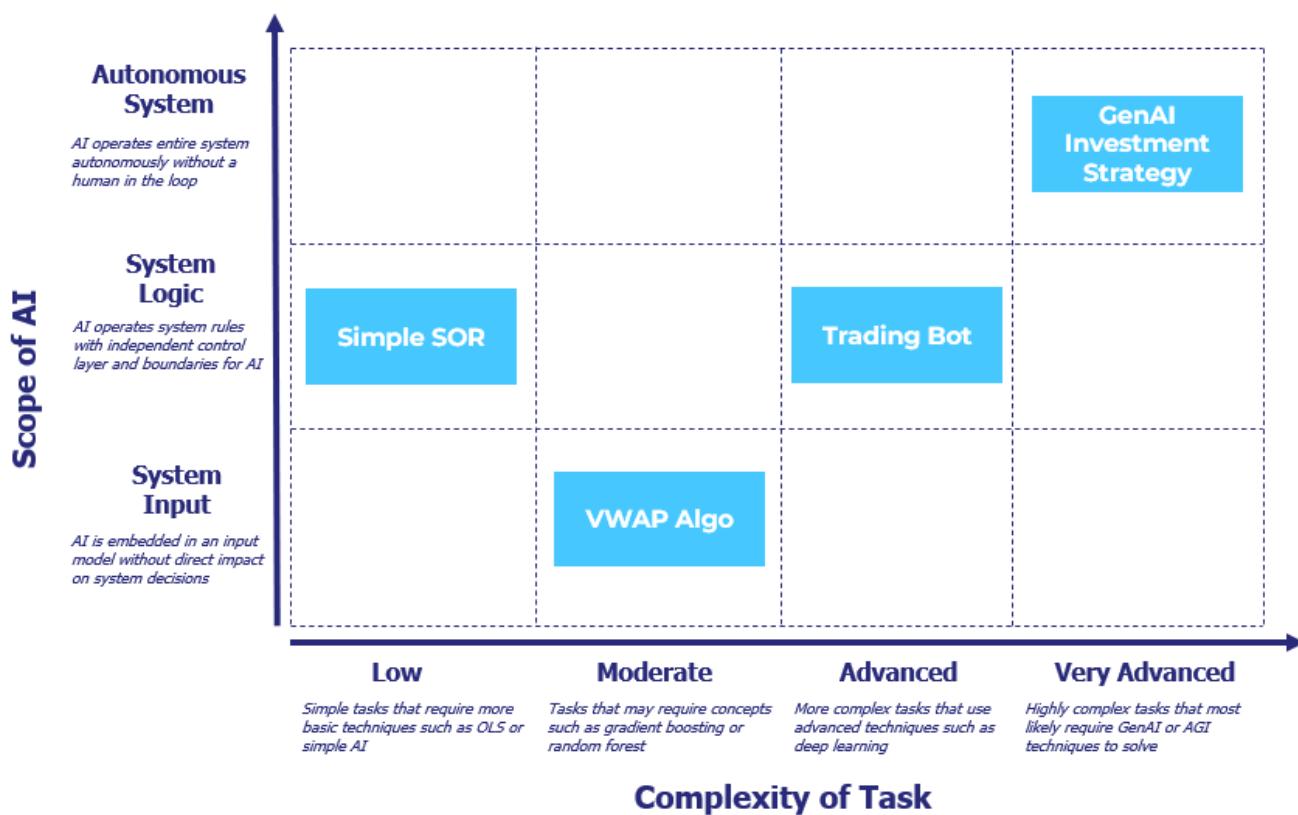


Figure 2 illustrates the spectrum of AI use cases, spanning limited applications where AI models are used as inputs without direct impact on system outputs, through to AI-driven system rules operating with human supervision or on a fully autonomous basis.

An example of a limited-scope AI use case could be a volume-weighted average price (VWAP) execution algorithm that incorporates AI-based price and volume models as inputs, while performing order splitting, scheduling and execution without using any further AI logic. In this configuration, AI informs the trading system but does not directly determine system actions or outputs. By contrast, examples of AI-driven system logic include a SOR that recommends or selects trading venues for execution. The decision taken by a SOR, such as choosing between

venues, may not be inherently complex but can be executed at scale and speed. A trading bot that seeks to replicate human judgement is a more complex task, even if it operates at a lower scale or frequency. In the longer term, autonomous investment strategies may provide AI models with broader discretion to invest, trade and monitor performance, potentially making use of generative AI techniques.<sup>1</sup> In this case, the entire rules-based code would be replaced with AI rather than individual components of the system.

Where AI is used to perform statistical modelling that generates inputs into a trading system without further influencing system outputs, it performs a role comparable to that of conventional econometric models. In such cases, the model estimates variables with a degree of uncertainty, as illustrated in the example above of the use of AI models in an execution algorithm to predict market prices and volumes. While AI introduces some additional considerations, the associated risks are largely consistent with those addressed by conventional model risk management frameworks (see Table 2).

**Table 2: Model risk management challenges with AI considerations**

Category	Issue	AI-specific risk considerations
<b>Data Management</b>	<ul style="list-style-type: none"> <li>Insufficient quality of input data</li> <li>Lack of transparency of training data</li> <li>Data ownership challenges</li> </ul>	AI models typically have a high reliance on existing and new data.
<b>Model Performance</b>	<ul style="list-style-type: none"> <li>Incorrect or inappropriate decision-making</li> <li>Scarcity of AI modelling expertise</li> </ul>	Model complexity or lack of research expertise could hinder the detection of model underperformance.
<b>Model Governance</b>	<ul style="list-style-type: none"> <li>Lack of accountability for AI models</li> <li>Scarcity of AI risk management expertise</li> </ul>	Model risk management guidelines and validation may need to be enhanced for AI model requirements
<b>Operational resilience</b>	<ul style="list-style-type: none"> <li>Failure of critical model infrastructure</li> <li>Susceptibility to cyber threats and fraud</li> <li>Insufficient technical expertise</li> </ul>	AI models may have different single points of failure or dependencies.

Regulators have established expectations for how firms identify, manage, and govern model risk within their trading infrastructure. These frameworks include the *Supervisory Guidance on Model Risk Management* (SR 11-7), issued jointly by the Office of the Comptroller of the Currency and the Board of Governors of the Federal Reserve System, as well as more recent UK guidance such as the Prudential Regulation Authority's *Supervisory Statement (SS1/23) on Model Risk Management Principles for Banks*. These frameworks are principles-based and technology-agnostic and therefore apply equally to AI models and to more traditional modelling approaches.

As modelling techniques evolve and advances in technology and computing power continue, it remains important that model risk management frameworks are implemented in a manner that keeps pace with these developments. This may require adjustments to model risk

<sup>1</sup> Additional case studies are set out in the annex.

management guidance, tools, or practices to address specific characteristics of AI models, such as increased scale or reduced interpretability of intermediary results. FMSB has previously published a *Statement of Good Practice on the Application of Model Risk Management Frameworks to Electronic Trading Algorithms* and similar good practices may be appropriate for the application of model risk management frameworks to AI models. Notwithstanding this, the underlying principles of existing model risk guidance provide a robust foundation for the governance of AI models, and do not, at present, indicate a clear need for a separate, AI-specific regulatory framework.

## Trading risk

Market-facing AI models are integrated into electronic trading systems that are subject to regulatory requirements governing algorithmic trading. These requirements mandate the implementation of appropriate controls across algorithmic trading infrastructure. Such controls are preventive and risk-limiting in nature and include, for example, price collars, maximum order size limits and position or risk limits linked to P&L exposure. Importantly, these controls sit outside the AI layer and operate as independent gateway controls before any model output interacts with the market. These controls mitigate AI model risk by constraining the maximum risk exposure that a model can assume. Where AI model behaviour is not adequately bounded by such gateway controls, further consideration may be warranted to ensure that the overall control environment remains appropriate, including the use of additional detective controls capable of identifying unexpected outcomes or errors in a timely manner.

While the core principles of model risk management and algorithmic trading control frameworks are designed to accommodate a wide range of stochastic models, the complex nature of advanced AI techniques and their novel deployments present a more pronounced challenge. When AI models operate with less constrained or more generalised objective functions than traditional modelling approaches, they may produce outputs where the rationale is difficult to interpret. As the novelty of application of these techniques continues to evolve, intermediary results may become unavailable, and it may not always be possible to trace back model outputs to specific input factors.

Confidence in the use of a model relies on robust validation, performance monitoring and testing of model outputs. These requirements apply equally to AI and non-AI models. Model validation should establish, at a minimum, that outputs are reproducible under equivalent conditions. Ongoing performance monitoring provides an additional layer of assurance by enabling firms to assess whether model outputs remain accurate and fit for purpose over time. Even where the specific drivers of an output cannot be readily explained, outcomes can be evaluated retrospectively. For example, while there may be challenges attributing the contribution of different factors to a particular price forecast supplied to an execution algorithm, the accuracy of that forecast can be observed and assessed after the event.

As technology and AI use cases evolve, specific guidelines, metrics and validation practices may need to adapt and market participants should continue to assess the adequacy of their control frameworks. Notwithstanding this, the core principles of model risk management remain applicable to AI models. Algorithmic trading controls provide an additional perimeter of real-time, independent safeguards around model outputs, enabling swift detection of errors and acting as effective risk-limiting guardrails. These controls are likely to be particularly important for more complex AI models, where increased complexity may further limit transparency. A sustained focus on monitoring model outputs and maintaining robust real-time controls provides a resilient and forward-looking approach to governing AI models as their use continues to expand.

## Systemic risks

Model risk management and algorithmic trading control frameworks enable market participants to manage AI model risks within well-established governance structures. Together, they provide an independent layer of controls between AI models and the market, limiting the direct impact of model behaviour and constraining potential risks to the wider ecosystem. This would change in circumstances where an AI system operates with full autonomy, without human supervision or intervention, or where AI is embedded within the control layer itself. While such configurations are not currently observed in wholesale markets, they are conceptually possible and could give rise to new or heightened systemic risks. Fully autonomous AI systems (as are beginning to emerge in other industries with different risk profiles) could, for example, influence price formation, affect supply and demand dynamics, or alter the behaviour of other market participants.

At present, the potential for AI models to create systemic risks for the ecosystem appears limited, not least because they are not sufficiently autonomous and do not operate outside of the confines of human supervision and intervention. Many of these potential risks are forward-looking and remain uncertain in terms of likelihood and impact:

- Synchronised AI-driven trading behaviour could amplify price movements and short-term volatility. This risk may be more pronounced where market-facing AI systems determine investment decisions, such as in short-horizon trading strategies.
- Coordinated or correlated behaviour among AI systems, which with sufficient autonomy and limited independent supervision, could adversely affect market fairness and efficiency. Such outcomes would likely depend on the design of objective functions that incentivised such behaviour.
- Failure or disruption of shared AI infrastructure could affect multiple market participants simultaneously. This risk is more closely linked to technological dependencies and infrastructure rather than AI modelling techniques themselves.
- Increased industry concentration or the emergence of new market entrants in the AI value chain reflect similar dynamics observed in other technology sectors. Such developments would be driven primarily by broader technology and market structure factors, rather than by the behaviour of individual AI systems or models.

Existing trading technology already presents potential systemic risks, as evidenced in various flash crashes and similar events. These risks stem from potential technology failures and are amplified when they occur in highly scaled processes or during highly concentrated periods of market activity such as benchmark points. While not a result of AI, these risks are still present in market-facing AI use cases due to its heavy reliance on technology.

For the purposes of long-term horizon scanning, Table 3 highlights a number of potential systemic risks that could arise from market-facing AI systems where the level of autonomy is increased in a way that is not observed today. Reduced transparency associated with some AI techniques could slow the detection, escalation and containment of such risks.

**Table 3: Potential systemic risk considerations of AI systems**

Category	Issue	Potential systemic risk
<b>Model commonality</b>	<ul style="list-style-type: none"> <li>AI systems which have been developed using similar techniques and/or training data allow flaws or biases to spread</li> <li>Correlated AI strategies create crowding and herding</li> <li>Learning techniques synchronise AI system actions</li> </ul>	Potential for exacerbated price reactions, increased volatility or destabilised markets
<b>AI model coordination</b>	<ul style="list-style-type: none"> <li>AI systems converge to mutually beneficial pricing</li> <li>Reciprocal adaptation reduces competition</li> <li>AI-to-AI system interaction increases the potential for coordination</li> </ul>	Potential reduction of market fairness and efficiency
<b>Collective infrastructure</b>	<ul style="list-style-type: none"> <li>Overreliance on AI systems across the ecosystem</li> <li>Shared AI system vulnerabilities</li> </ul>	Potential for market-wide outages that endanger market functioning and resiliency
<b>Competition dynamics</b>	<ul style="list-style-type: none"> <li>Concentration of data or third-party AI infrastructure in a few actors could endanger competition</li> <li>New market participants may emerge outside of existing regulatory framework</li> </ul>	Potential for concentration in the market for data and AI infrastructure, undermining access, fair pricing and competition and giving the owners of AI models a systemically important role

## Accountability

As AI models evolve and their use in markets expands, a greater share of trading-related tasks will be performed by machines rather than by individuals. This shift alters how decisions are taken and where responsibility resides, with greater reliance placed on system design, implementation and oversight. As a result, considerations of conduct risk and accountability increasingly extend beyond individual trading decisions to the governance of AI models and systems.

Clear accountability remains essential as AI capabilities develop. AI models and systems should have clearly defined human owners, analogous to the accountability frameworks applied to traders and trading management today. Where AI systems exhibit adaptive behaviour or operate with a higher degree of autonomy, accountability for system actions and outcomes should continue to rest with humans. Effective accountability requires that responsible roles are equipped with appropriate tools and capabilities, including monitoring metrics and alerts for system outputs, manual and automated intervention mechanisms, and management information sufficient to support oversight and timely action.

Structural change and technological evolution are longstanding features of wholesale markets. Market-facing AI models and systems have the potential to enhance productivity and market efficiency. When combined with robust governance, controls, and continued human accountability, their use can support fair and effective markets as these technologies continue to develop.

## 4. Assessing AI in trading systems

### Contextual considerations

This section brings together the key contextual considerations for firms when evaluating AI use cases within trading systems and applies them to a set of illustrative case studies. Together, these broader questions and more detailed checkpoints draw on the observations from earlier sections and provide a practical starting point for contextualising AI use cases within electronic trading systems. The considerations are illustrative and non-exhaustive.

**Figure 3: Contextual considerations for AI use cases**

#### **Step 1 – What is the purpose of the use case and how does it work?**

- Is the use case market-facing?
- How complex is the task that AI is performing?
- What is the scope of AI: input module or system logic?
- How much autonomy and self-learning does the use case exhibit?
- Does the use case involve human supervision and intervention?

#### **Step 2 – What are the key risks associated with the use case?**

- What undesirable outcomes may occur?
- Does AI significantly alter the use case's risk profile?
- How do the risk characteristics of AI models compare with those of traditional models?
- How material would errors be for the firm and wider stakeholders?

#### **Step 3 – What controls are needed on the model inputs, outputs and outcomes?**

- Can AI model outputs be reproduced for validation?
- Is the control framework suitable for the scale and complexity of the task?
- Do existing MRM guidelines and metrics adequately assess the risks associated with the AI model in the relevant use case?
- Are algo trading controls sufficiently risk-limiting for the AI use case?
- Are the controls capable of detecting model failures in a timely manner?
- What skills do individuals running the system require?

## Case study 1 – Electronic trading system with AI-powered components

<b>Synopsis</b>	<b>Liquidity provider operates a complex electronic trading system with multiple AI-driven components</b>
<b>Relevant risk categories</b>	Model risk, trading risk
<b>Key contextual considerations</b>	<ul style="list-style-type: none"> <li>Identify AI components in a complex trading system</li> <li>Assess market-facing nature and complexity of AI analytics</li> <li>Apply appropriate output-based controls</li> </ul>
<b>Main observations</b>	<ul style="list-style-type: none"> <li>Decision-making level and market-facing nature of AI are limited</li> <li>AI introduces limited incremental risks to the electronic trading system compared to non-AI models</li> <li>Liquidity provider already captures the risks in their existing model risk management and electronic trading control frameworks</li> </ul>

### Scenario

A liquidity provider streams two-way prices directly to clients and venues in spot FX through its existing electronic trading infrastructure. The liquidity provider identifies and implements some enhancements with the help of different AI techniques and at different points in its electronic trading workflow.

<b>Step 1:</b>	The electronic trading system continuously takes in a combination of structured and unstructured data to support pricing and execution. This includes updated market data, internal position and exposure data and data from unstructured sources like news feeds, central bank data and analyst reports. The unstructured data gets processed with the help of LLMs to produce new sentiment and macro signals. Due to latency, the use of LLM output is kept several steps removed from direct interaction with the client or venue and limited to signal generation. As the output is distilled into a signal, this can be qualitatively evaluated using long-established techniques.
<b>Step 2:</b>	Once the electronic trading system has received the inputs from step 1, it processes them through a second layer of analytical models. Alpha generation models process the signals to predict prices or client interest. Risk models evaluate internal positions against the liquidity provider's risk appetite. Market impact models estimate expected transaction costs of trades. All models are updated and refined based on live market data, position data and execution data and can be qualitatively evaluated using established techniques.
<b>Step 3:</b>	The outputs of the second layer of models are then fed into execution and pricing algorithms. The electronic trading algorithm determines how, where and when to execute trades. Reinforcement learning and adaptive algorithms are used to optimise order execution and dynamically skew client pricing based on real-time market conditions. The performance of these algorithms can be measured using simulation techniques common with algorithmic trading.

**Step 4:** Prior to trades being executed, the electronic trading system applies automated controls to ensure that all trades adhere to position, size and price limits as well as further gateway controls.

### Potential risks

The electronic trading system consists of multiple analytics components and sequential layers. Most components, such as the risk model or the market impact model, are based on standard econometric modelling techniques. Some models, such as the signals, have been enhanced to use AI techniques such as natural language processing. All models are isolated input components that do not directly trigger market-facing actions. The types of risks of these models are common statistical modelling challenges (and not exclusive to AI).

In step 2, all models get updated based on live data. This implementation of the model should undergo the same model risk management procedures and implementation testing as any other model updating in real-time. It also requires real-time detective controls or guardrails commensurate with its risk.

The trading system consumes all AI model outputs into its engine and updates the system outputs based on real-time data. Additionally, within the execution and pricing algorithms, the liquidity provider uses reinforcement learning techniques for optimising pricing and execution decisions. They are being applied closer to the market-facing processes, even if not directly interacting with the market. This heightens model risk, similar to the existing models in the trading system.

### Controls and potential impacts

The signal and modelling steps are governed by MRM frameworks that ensure model soundness, monitor model performance, and provide controls over input data. The use of AI techniques increases the importance of rigorous data management, a risk that is addressed within existing MRM frameworks.

The reinforcement learning model in step 3 is more complex. It produces recommendations that are designed to get actioned in the market. Before execution, the electronic trading system applies risk-limiting controls to ensure that all trades adhere to restrictions, capping the risk of the AI model output. All gateway controls are independent and outside the AI layer.

Overall, the use of AI has in some instances changed the risk profile of the use case, heightening some risks of the electronic trading system. However, the risks are comprehensively managed in robust independent model risk management frameworks and trading risk control layers.

## Case study 2 – Market-making agents

Synopsis	AI-driven strategies directly adapt behaviour to market conditions
<b>Relevant risk categories</b>	Model risk, trading risk, conduct risk
<b>Key contextual considerations</b>	<ul style="list-style-type: none"> <li>Identify level of autonomy of AI use case</li> <li>Assess risks from directly market-facing nature of use case</li> <li>Understand whether AI alters the risk profile</li> </ul>
<b>Main observation</b>	<ul style="list-style-type: none"> <li>The use case is market-facing with a high potential for an adverse impact on market outcomes</li> <li>The types of risk, while high, are the same as for highly automated non-AI strategies</li> <li>Need for independent controls that can detect conduct patterns and need to align AI use case with humans accountable for oversight</li> </ul>

### Scenario

Two large liquidity providers independently deploy reinforcement learning (RL) agents to support their electronic market-making in mid-cap corporate bonds. These agents are designed to optimise trading performance by continuously learning from their own execution outcomes and from evolving live market conditions. Each agent independently adjusts bid-ask spreads, order sizes, and quotes with limited human input. The agents operate in parallel across overlapping venues and time frames, adapting strategies in response to real-time feedback.

### Potential risks

Multiple reinforcement learning agents may begin reacting and adjusting behaviour to the same market signals or events in near real-time. If their strategies are similarly structured, they may withdraw or adjust liquidity simultaneously, unintentionally amplifying volatility and triggering sharp, short-lived price dislocations.

An RL agent interprets short-term price movements as trading signals, not recognising that its own activity is driving the change. It accelerates execution in response, further pushing the market and reinforcing the behaviour, leading to procyclical distortions.

Conceptually, and even if not designed to do so, an agent trained to maximise execution quality may discover that placing and then quickly cancelling large orders on the opposite side of the book influences fill probability. Without constraints or controls, this behaviour could emerge as an unintended strategy.

## Controls and potential impacts

The use case involves AI-driven strategies that, with limited human intervention, may potentially produce outcomes that are highly correlated with other market participants. If the AI model is poorly constructed, it may additionally underperform.

These types of correlation challenges are the same challenges that any highly automated trading system, or a human trader at lower scale, face and that MRM frameworks address in their model validation, performance monitoring and testing pillars.

The high level of independent actions without human inputs makes behaviours such as unintended spoofing or other forms of market manipulation conceptually possible. This underlines the importance of an independent control layer with detective controls that identify and rule out prohibited behaviour patterns.

The more autonomy an AI use case exhibits, the more important the independence of controls become. This level of autonomy of AI would create challenges in the existing regulatory framework as to how a system is tested and whether the control framework is reasonably designed.

While it may appear as if collusion becomes more likely with AI agents, it is in principle not different from the collusion risk of two human traders or two firms with highly automated strategies that process market data.

As AI models replace human actions in this case, it is critical that there still is clear accountability for each AI agent: a clear 'owner' of the code and a clear 'business owner' who oversees the usage and risk of a market-facing AI strategy. This should be no different from existing trading systems that are highly automated.

### Case study 3 – AI-based Customer Relationship Management (CRM) for client pricing

Synopsis	Firm applies AI-driven client data to its tiered pricing engine
<b>Relevant risk categories</b>	Conduct risk, operational risk
<b>Key contextual considerations</b>	<ul style="list-style-type: none"> <li>Identify sensitivity and heightened risk that arises from input data</li> <li>Assess how material an error could be for stakeholders</li> <li>Ensure that level of controls is sufficient for identifying bias</li> </ul>
<b>Main observation</b>	<ul style="list-style-type: none"> <li>The firm's AI activity is market-facing but with limited systemic risk</li> <li>Control frameworks need to have the ability to recognise where potential bias can harm institutional clients</li> <li>Staff need sufficient expertise for applying AI to sensitive data</li> </ul>

#### Scenario

A firm deploys a retrieval-augmented generation (RAG) system to support client-specific pricing decisions in its electronic request-for-quote (RFQ) workflows. The tool is embedded within the firm's client management infrastructure and is designed to flag relevant servicing and behavioural insights when a pricing request is received. When a client sends a pricing request, the AI pricing model scans internal client relationship management notes, including emails and phone calls with clients, servicing records and historical trade data. Insights retrieved from the CRM system are returned as a short prompt to the sales trader, or, in some cases, passed directly to the pricing engine to inform quote calibration. For example, the pricing model highlights that a client has significantly reduced flow over the last two months and has not met prior agreed targets. This triggers a recommendation to downgrade the tiering of the client for pricing purposes, adjusting the spread before the quote is finalised.

#### Potential risks

Unlike more traditional techniques, if challenged, it may not be clear why the pricing model triggered a pricing change. This could make it harder to explain pricing decisions and demonstrate fairness to clients. The output of the pricing model could result in different pricing outcomes for similar clients across regions that are hard to explain. It may also be harder to identify when a model has an in-built bias.

Additionally, in this case study the human trader may begin to rely on AI-generated summaries without applying appropriate judgment or commercial awareness – in the same way that a trader may over-rely on non-AI aids as well.

### Controls and potential impacts

MRM frameworks seek to ensure that models are bias-free. In this AI use case, it is important to ensure that the guidelines contain testing requirements or model validation requirements that help reduce the risk of bias as much as possible, in particular when model results are hard to understand and market participants cannot rely on interpretability for assessing model bias.

Staff who rely on these AI models when dealing with customers need to be skilled enough to understand a model's intended use case and its limitations. Responsibility for flaws or erroneous decisions cannot lie solely with a model.

## 5. Conclusion

AI is becoming an increasing feature in financial markets. Market participants continue to face the challenge of balancing the innovation potential of AI with the need for robust risk management and appropriate safeguards. In trading, AI is typically deployed as part of larger electronic trading systems that are already highly automated, scalable, and analytics-driven. In most current use cases, AI models operate within defined and contained components of these systems, providing inputs rather than directly interacting with the market, and remain subject to independent control layers and human supervision. In that context, AI is not novel but is an enhancement of existing statistical inference techniques. While future developments may enable AI systems to exhibit greater autonomy and adaptive behaviour, such configurations are not currently observed in market-facing applications.

Nevertheless, AI is already supporting use cases that address increasingly complex tasks and objective functions. The risks associated with these use cases arise primarily from the nature and complexity of the activities performed, rather than from the specific AI techniques employed. Assessing AI-related risks in the context of the relevant use case and AI's role within the trading system therefore remain critical. Many such risks are already addressed through established control frameworks, including model risk management and algorithmic trading requirements. However, the scale, speed, or novelty of certain AI applications may place pressure on existing guardrails and give rise to additional considerations to which market participants should remain alert.

## Annex 1 - AI definitions

Notable existing definitions of AI include:

Source	Definition
<b>Organisation for Economic Co-operation and Development, OECD Recommendation of the Council on Artificial Intelligence, adopted 22 May 2019, amended 8 November 2023</b>	<p><i>'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.'</i></p>
<b>National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework, NIST AI 100-1, January 2023</b>	<p><i>'An AI system is an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.'</i></p>
<b>Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Article 3(1)</b>	<p><i>'AI system' means 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;</i></p> <p><i>Recitals and interpretative guidance relating to Article 3(1), further clarify the scope of what constitutes an AI system in a market-facing context.</i></p> <p><i>Recital 12 states that: 'the definition should be based on key characteristics of AI systems that distinguish it from simpler traditional software systems or programming approaches and should not cover systems that are based on the rules defined solely by natural persons to automatically execute operations.'</i></p>
<b>Commission Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (AI Act)</b>	<p>Recitals 41 and 42 also provide that:</p> <p><i>'Some systems have the capacity to infer in a narrow manner but may nevertheless fall outside of the scope of the AI system definition because of their limited capacity to analyse patterns and adjust autonomously their output. Such systems may include:</i></p> <p><i>Systems used to improve mathematical optimisation or to accelerate and approximate traditional, well established optimisation methods, such as linear or logistic regression methods, fall outside the scope of the AI system definition. This is because, while those models have the capacity to infer, they do not transcend 'basic data processing'. An indication that a system does not</i></p>

*transcend basic data processing could be that it has been used in consolidated manner for many years. This includes, for example, machine learning-based models that approximate functions or parameters in optimization problems while maintaining performance. The systems aim to improve the efficiency of optimisation algorithms used in computational problems. For example, they help to speed up optimisation tasks by providing learned approximations, heuristics, or search strategies.'*

## Annex 2 – Bibliography

- **European Union**, Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).
- **European Commission**, Commission Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (Artificial Intelligence Act).
- **European Commission**, Commission Delegated Regulation (EU) 2017/589 supplementing Directive 2014/65/EU with regard to regulatory technical standards on organisational requirements of investment firms engaged in algorithmic trading (RTS 6).
- **Financial Markets Standards Board (FMSB)**, Statement of Good Practice for the Application of a Model Risk Management Framework to Electronic Trading Algorithms, April 2024.
- **Prudential Regulation Authority (Bank of England)**, Model Risk Management Principles for Banks, Supervisory Statement SS1/23.
- **Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency**, Supervisory Guidance on Model Risk Management, SR 11-7.
- **Securities and Exchange Commission**, Risk Management Controls for Brokers or Dealers with Market Access, 17 CFR § 240.15c3-5.
- **Hong Kong Securities and Futures Commission**, Code of Conduct for Persons Licensed by or Registered with the Securities and Futures Commission, Chapter 18, as amended.