



Governing Autonomous Intelligence in International Financial Systems

Yannis Ioannidis
Asst Professor

Abstract

Autonomous machine learning — defined here as a system's use of its own outputs to drive further decision-making without human intervention — has had a profound impact on global finance. This influence is evident in cross-border flows of funds for transactional business in products and services, financial exchanges, and financial assets. As these flows grow more automated, a clearer separation has emerged between the actual originator of a transaction and the funds-handling entities in other jurisdictions that provide settlement or routing services. This separation introduces risk and compliance challenges, since funds-handling entities must deploy systems capable of enforcing Know Your Customer, Anti-Money Laundering, sanctions, and other applicable regulatory requirements across borders.

Translating cross-border data-sharing needs into practice, however, remains difficult. For the data-sharing requirements that arise in cross-border financial interactions involving autonomous machine learning, a risk-based allocation of roles has motivated efforts to develop data-sharing agreements tailored to specific jurisdictions. The underlying principle is to route transactions through jurisdictions with lower-risk data-regulatory profiles, enabling near real-time information flow while minimizing the time sensitive data must reside in jurisdictions with more stringent regulatory requirements.

Keywords: Autonomous Machine Learning, Cross Border Financial Transactions, Real Time Decision Systems, Funds Flow Automation, Transaction Settlement Networks, Global Payment Infrastructure, Know Your Customer Compliance, Anti Money Laundering Controls, Sanctions Screening Systems, Risk Based Role Allocation, Cross Jurisdiction Data Sharing, Data Regulatory Interoperability, Financial Crime Prevention, Autonomous Compliance Systems, Jurisdictional Risk Management, Real Time Information Exchange, Data Residency Constraints, Regulatory Technology Integration, Global Financial Governance, Machine Driven Risk Assessment.

1. Introduction

The governance and compliance of cross-border finance involve interactions and transactions with foreign entities through the financial system and involve a broad range of activities such as real estate purchases, international investments, money transfers, and tourism. These financial flows typically involve data sharing between jurisdictions, making them subject to data governance-related laws and regulations affecting the physical flow of data and the infrastructure of communications used.

The advent of autonomous machine learning creates new possibilities for cross-border uses, supports decision-

making, risk assessment, and support provisioning, automates compliance with KYC and AML laws, authenticates business scenarios, routes payments around congested corridors, and issues central bank digital currencies bridging sovereign currencies in real-time. Several transactions have already been taken to market. However, the seamless automation of cross-border transactions remains some distance off. Key gaps that are delaying the final destination are deficient regulatory alignment, the lack of political agreement for mutual recognition of compliance, and the ability for regulators, especially in non-aligned jurisdictions, to monitor risk and have redress.

1.1. Overview of Autonomous Machine Learning's Impact on Global Financial Interactions

Autonomous machine learning (ML) fundamentally reshapes cross-border finance by refining transaction channels and settlement processes, focusing on verification of transaction parties and counterparty creditworthiness, and optimizing decision-making and risk-assessment tasks that underpin transactions or support supervisory analysis. Specified risk classification criteria enable delegated functions to rely increasingly on autonomous decisions, thereby transferring costs to users or generating profits while improving execution quality. Corporates, institutions, and retail clients also benefit, although profit-driven diversity could fuel credit and liquidity risk.

Decision-making processes increasingly rely not only on the draft responses produced by autonomous ML agents but also on the quality, accuracy, reliability, and integrity of the underlying data sources. Transparency concerns focus not only on the entities assigned the task but also on the input data sources rather than the ML execution's decision-making itself. Enhanced resilience from ML-supported fraud detection and prevention systems complements the associated profitability but remains critical in the absence of profit motivation. Gaps and omissions in global cross-border business and policy-making activities are expected to be addressed through autonomous ML-supported services, with business incentives driving safe implementation. Compliance with formal writing requirements and all regulatory and policy aspects delegated to autonomous systems with specialized expertise and knowledge facilitates supporting tasks for all actors. Business decision-making in complex data and operational environments also increasingly relies on generative AI agents running in voice format.

Autonomous ML analytics performs continuous danger mitigation. Businesses, especially multinationals, have shifted towards a mined-data approach, governed mainly by data sovereignty, but it still requires proper credit risk, capital, or liquidity assessment decisions. Considering ML's impact on cross-border KYC/AML processes or

international sanctions compliance that do not imply capital transfer, industry-enabled incentives and public-private partnerships support operational intelligence scratchpads and training environments for testing non-production systems, even if business-sensitive. In turn, market settlements with data and agreement commensurate coding complexity and audibility are continuously optimized using these same solutions. Transaction parties can autonomously code KYC information in markup languages associated with certificate payloads in smart contracts or executable agreements hosted on transaction ledgers. Probability coding exposes verification paths and necessary capabilities, and local drift and update training through economic or voluntary agreements relies mainly on local data classification quality.



NECOINETS PEESAND FOR CROSRSSEMENTE AI RETIR AL CONTROINS ETHI AOL RETAILLIENTS


Fig 1: The Autonomous Frontier of Global Finance: Algorithmic KYC, Smart Contract Integration, and Risk Mitigation in Cross-Border Settlement Ecosystems

2. The Landscape of Cross-Border Finance in the Age of Autonomous Machine Learning

The current and future landscape of cross-border finance in the age of autonomous machine learning evolves through the lenses of transaction channels, settlement, and

regulatory reporting. On the one hand, major financial markets converge through the deployment of neural networks and other learning algorithms, fostering new automation and efficiency levels; on the other, a fragmented regulatory approach hampers the naturally borderless nature of the internet.

Data sovereignty and interoperability are key considerations. Transactional data belong to the jurisdiction in which the transaction is executed. Transparency—data storage and access protocols, processing methods, and authorizations for data access—is thus crucial to comply with the desire of governments, organizations, and people to protect their data and the datarights of their citizens. The evolution of cross-border transactions is propelled by machine learning invention and technology. Key elements include the automation of Know Your Customer (KYC) and Anti-Money Laundering (AML) controls; automated sanctions screening; the reduction of payment or settlement times to real time; and intelligent routing of transactions based on cost, time, and liquidity considerations. As efficiency increases across cross-border transaction dimensions, so do potential risks.



Equation A: Risk-based explainability requirement (Accountability/Transparency/Explainability)

A1) Define variables

- Let $r \in [0,1]$ be **perceived risk** of an activity (0 = low, 1 = high).
- Let $e \in [0,1]$ be the **required explainability/documentation level**.

A2) A minimal monotonic model

We want e to:

- be small when r is small,
- rise quickly around “medium risk”,
- saturate near 1 at high risk.

A standard choice is a **sigmoid**:

$$e(r) = \frac{1}{1 + \exp(-k(r - r_0))}$$

2.1. The Evolution of Cross-Border Financial Transactions Driven by Machine Learning

The evolution of cross-border financial transactions driven by machine learning (ML) can be traced through four phases: (1) intelligent assistance for customer onboarding and regulatory validation, (2) autonomous fulfilment of "Know Your Customer" (KYC) and "Anti-Money Laundering" (AML) regulatory requirements, (3) real-time settlement of the transaction itself, and (4) intelligent routing of capital flows to optimize transaction efficiencies between jurisdictions.

The first phase saw the inception of intelligent assistants powered by ML technology and Natural Language Processing (NLP), which vastly improved user experience and mitigated the resources needed to comply with regulatory requirements. M-Learning enabled chatbots and virtual agents capable of conversing in a natural manner with clients, generating trust in the proposed solutions. As such systems became commonplace, the need for remote yet rigorous KYC processes, and underlays enabling cryptographically secure identity validation, led to the



introduction of ChatGPT-like experience-powered services for KYC. Clients were thus able to perform the full KYC process naturally in a virtual environment, aided by an NLP assistant on the front-end, and having the requested documents automatically validated and analyzed by ML engines at the back. KYC and other onboarding requirements therefore became lighter than ever, thus greatly reducing latency. Currently, the focus is placed on achieving ML-driven automated KYC and AML through federated learning, Multi-Party Computation (MPC), or other consent-based methods.

3. Governance Frameworks for Autonomous Financial Systems

Governance frameworks for autonomous financial systems should aspire to articulate broad objectives, delineate key actors, and establish fundamental principles. The guidelines, allowances, and prohibitions ultimately shaping the day-to-day operations of such systems should stem naturally from an amalgamation of international standard-setting and national policy and regulatory choices. While the need for regulators to retain a risk-based approach—a focus on the riskiness of institutions, market participants, activities, and products, rather than attempting to regulate the deployment of specific technologies—has never been more universal, a shift toward a more outcome-based paradigm also seems advisable. This shift calls for the design of compliance regimes that express expectations in terms of outcomes that must be delivered by the financial ecosystem rather than strings of activities that must be performed. Action-specific rules can then apply only where the nature of the proposed activity or system gives rise to heightened confidence that activity-agnostic outcomes will not be achieved.

A more coherent approach to oversight becomes even more critical when the system in question is one designed to be operated by only partially trusted participants or by participants that are unknown to one another. It is their lack of trusted credentials that affords the participants in an

autonomous financial ecosystem the freedom to address actors and find services previously inaccessible, and it is their anonymity that enables unsanctioned participants to offer services without getting caught. No one would allow an aged uncle at risk of dementia to write instructions for complex brain surgery; similarly, no one should rely on financial institutions whose very ability to continue functioning depends on key people not having dementia. And yet the success of the cross-border financial ecosystem during the twentieth century relied on a large share of its various components being supplied by such participants or by otherwise equally unknown counterparties. The timeframe for autonomous financial systems adds the new and challenging requirement of not merely enabling undesirable transactions to happen quickly but doing so efficiently and silently while staying one step ahead of fraudsters with dubious moral compasses.

Symbol	Meaning
x	Feature vector describing party/transaction (KYC/AML inputs)
$s(x)$	Model score = $P(\text{illicit} \mid x)$
τ	Decision threshold for flagging / blocking
FP, FN	False-positive and false-negative rates
C_{FP}, C_{FN}	Unit costs of FP and FN outcomes
$j \in J$	Jurisdiction index on a cross-border route
c_j, ℓ_j	Processing cost and latency in jurisdiction j
ρ_j	Regulatory/data stringency risk score for jurisdiction j
p	A route (path) across jurisdictions
w_c, w_ℓ, w_ρ	Weights for cost, latency, regulatory risk in routing objective
Q	Net signed order flow (buy +, sell −)
ΔP	Price impact caused by flow
S	Quoted bid–ask spread
Π	Expected market-making profit (PnL)

3.1. Policy and Regulatory Alignment Across Jurisdictions

The increasingly autonomous nature of today’s financial systems—intelligence and task execution without human input—disrupts established patterns of policy and regulatory governance for cross-border transactions. These transactions are also known as cross-border financial exchanges or business travel; here, the Anglo-American concept of business travel is preferred because other borders are less important than the business value involved. Much of the previous analyses supports a different concept of policy and regulatory alignment. Autonomous machine learning (ML) changes aspects of decision-making, task execution, and information retrieval that affect the nature and direction of travel between financial jurisdictions. Coherent governance, compliance, and operational verification frameworks across jurisdictions become imperative, even for low-risk flows, if autonomous finance is to remain viable.




Fig 2: Incremental Harmonization: An Iterative Experimental Framework for Regulatory Alignment in Autonomous Cross-Border Finance

Disparate decisions at different times typically yield formal treaties and executive agreements, yet these result in more of a patchwork quilt than a coherent alignment strategy. Attempts to realign the patchwork typically take the opposite, guns-blazing approach, seeking to unify entire regulatory or compliance divisions helter-skelter across continents with a set of minor props, rather than simply allow cooperation and alignment to occur in a piecemeal fashion. Much smarter is to first identify a smallish set of jurisdictions that might be willing to experiment with their supervisory approaches, where obvious cooperation might yield sound benefits, then perform the experiment and see how it goes—so that the next set, now maybe even with a precedent to lean on—can easily align these new learnings with their own ideas and testing before actual implementation. After that, check what works and what does not and make changes as required and then proceed to the next set.

3.2. Accountability, Transparency, and Explainability

The decisions taken by autonomous machine learning (ML) do not require the same level of justification as for human action. However, from a compliance perspective, the stakeholder acting on the advice or undertaking the action should be able to ascertain the justification. The level of accountability, transparency, and explainability required depends on the perceived risk of the activity, degree of autonomy, and AI’s risk capabilities. Information needs can be deduced separately for advisory and autonomous AI decisions and actions.

Every AI decision made can be traced back to the data and models that influenced it, but the general public may lack the understanding to interpret that information. For society to have reasonable confidence that AI performs within its limits, organizations must communicate AI decisioning in simple, understandable terms, provide an assurance statement acceptable to the relevant jurisdiction, or undertake additional manual or enhanced monitoring activity. Key points include that delegated responsibility

systems should function as intended and that where AI plays, the traffic and fraud light is green. Formal documentation is required to meet the needs of regulators, internal audit functions, and external assurance providers, covering AI decision-making models to an appropriate depth based on risk-based and outcome-oriented principles. This documentation should also support model audit trails, and be available to those acting on model advice and those responsible for the action. Furthermore, for the supervisory activity to be successful and effective, records of model development activities, operating conditions, monitoring results and operational incidents related to reliance on or interfacing with the model should also be retained, so that they can be provided to the supervisory authority when required.

4. Compliance Mechanisms in Generative AI-Driven Financial Services


Compliance in a financial context generally refers to the capability and assurance that a business, department, or professional operator meets the applicable laws, regulations, guidelines, and specifications for its activities. Financial service firms operate across multiple jurisdictions that expect principals to fulfill all local supervisory and regulatory obligations and implement comprehensive compliance frameworks. Generative AI capability applications in financial services are often advisory in nature; however, advisory solutions could lead to potential autonomous actions in the future. Such actions are traditionally framed as the company’s responsibility, whether human or machine, hence risk management and control environments should include and evaluate regulatory and operational risks to fulfill compliance requirements.

An effective compliance mechanism environment encompasses the assessment, establishment, maintenance, and testing of procedures to enable compliance with relevant laws and regulations. The deployment of compliance-aware generative AI may not substitute a

controls framework but should affect them. Therefore, the management response must consider generative AI as an enabling technology allowing a company to fulfill compliance requirements.

For example, a bank trades across regulated markets, exchanges, and over-the-counter markets for its own account or on behalf of customers. Regulators expect the bank to put in place a control framework that prevents, detects, and mitigates trading failures and misconduct, including market abuse and manipulation. The trading operations rely and comply with the local rules of the jurisdiction the operations happen in and the rules that apply to the counterparties. Given the nature of generative AI, compliance is usually embedded in the solutions developed, especially for advisory services. The market conditions and the news information change with higher speed than the response that the system can provide. Therefore, when deploying generative AI responses, it is important to define the target for the advisory nature of the responses.

Figure B — Profit condition: spread must exceed mispricing losses (illustrative)



Equation B: KYC/AML decision threshold (false positives vs false negatives)

KYC/AML and sanctions controls and monitoring false positive/false negative rates.

B1) Score model

Let the ML system output:



$$s(x) = P(Y = 1 | x)$$

Where:

- $Y = 1$: illicit / suspicious
- $Y = 0$: legitimate
- x : transaction/customer features

B2) Threshold rule

Flag/block if:

$$s(x) \geq \tau$$

B3) Expected cost of a decision (step-by-step)

Two error types:

- **False Positive (FP):** block a good customer $\rightarrow$ cost C_{FP}
- **False Negative (FN):** allow illicit activity $\rightarrow$ cost C_{FN}

If we **block** (predict 1):

- expected cost $= C_{FP} \cdot P(Y = 0 | x)$

If we **allow** (predict 0):

- expected cost $= C_{FN} \cdot P(Y = 1 | x)$

So choose “block” when:

$$C_{FP}(1 - s(x)) \leq C_{FN}s(x)$$

Solve step by step:

1. Expand:

$$C_{FP} - C_{FP}s(x) \leq C_{FN}s(x)$$

2. Move the $s(x)$ terms to the right:

$$C_{FP} \leq (C_{FP} + C_{FN})s(x)$$

3. Divide:

$$s(x) \geq \frac{C_{FP}}{C_{FP} + C_{FN}}$$

4.1. Risk Management and Control Environments

In the context of generative AI-enabled systems – and more broadly in relation to an autonomous machine learning environment – the concept of a Risk Management and Control Environment must extend into a number of domains that may not traditionally have been associated with enterprise risk management. These domains include identification of new risk taxonomies, establishment of Control Activities associated with the new risk categories, Monitoring of the designated Control Activities, Identification of Incidents and associated response, and definition of Regulatory Reporting obligations.

The relatively recent emergence of generative AI technologies presents both a significant opportunity and risk for financial institutions – indeed, for all enterprises – and, as the hype recedes, it is increasingly evident that these technologies are more than just another bandwagon on which to leap. Generative AI represent a fundamental paradigm shift in business operations, have already demonstrated how ease of use can result in significant security and privacy risks, and are therefore a new and emerging risk category in their own right. Given the capability of generative AI to automate, support or augment a wide variety of decision-making and operational activities – including potential Take Action functions – it is highly unlikely that enterprises can simply contemplate their use without a robust risk management and control framework together with aligned Regulatory Reporting policy.

5. Technical Foundations of Autonomous Machine Learning in Finance

Autonomous machine-learning systems for financial service activities—such as transaction routing, trade execution, transaction monitoring, KYC, AML, and



sanctions screening—operate in a highly interconnected environment. Their design, operating model, and risk controls must address technical conditions and governance expectations established by the wider ecosystem in which they exist. An autonomous ML solution is likely to have been developed and is being operated in an environment that automates the entire model and risk management life cycle and has strict external cybersecurity and resilience requirements. Technically minded users should consult the following sections to understand how the architecture, data governance, and risk controls of an autonomous ML solution support the governance and compliance requirements of other services that it interacts with, either directly via calls or indirectly via an API.

A dedicated model-governance framework should also exist to support the model’s development, validation, deployment, monitoring, versioning, off-switch criteria, and retirement. These aspects should be documented and periodically revalidated. At a minimum, they should cover model performance during its active life cycle—track record on authorisation (false positive and false negative rates per classification category), trend analysis on declared sanctions, and any detected model drift as indicated by monitoring—and detail the reasons for any model-version rollbacks or other off-switch activity and how these actions are communicated to the relevant non-technical stakeholders. Model performance in the pre-deployment period should also be carefully controlled, with the retention of all test data and relevant communications post-validation for audit and compliance purposes.

Risk tier	Max autonomy (0–1)	Explainability depth	Monitoring frequency	Control mode
Low	0.2	Light	Monthly	Post-hoc sampling
Medium	0.5	Moderate	Weekly	Pre-trade

Risk tier	Max autonomy (0–1)	Explainability depth	Monitoring frequency	Control mode
				checks for sensitive corridors
High	0.8	Strong	Daily	Real-time gating + human escalation

5.1. Model Governance and Lifecycle Management

Governance and compliance connect towards the effective use of autonomous ML in financial institutions. These objectives naturally extend to AML activity. However, as described, capabilities extend the KYC and sanctions processes from guiding advisory to allowing operational-enabled services. The trackable decisioning reduces the Friction Risk of these elements, especially at scale. This lays a challenge for the Control Environment to enable, maintain, and demonstrate supervisory compliance in autonomous mode. The Governance Framework establishes expectations for the Test of Quality, Test of Control, and Test of Intelligence. Within these pillars, capability design and enabling processes—Model, Data, Application, Information Security—embody the associated tracking requirements of decisioning, model ownership, and KYC simplification.

Model Governance focuses on the model development lifecycle from design to production. Completion of the Test of Quality confirms the general-purpose validation for model enabled operations. Model Governance is primarily on the end model sections and incorporates Model Information Control requirements. The Data Governance Control Activity includes a layer above the model to enable

a risk assessment of external inputs. Criticality, reputation, regulatory, and operational risks all link back to Data Governance and must be assessed as part of the approval process for the ML solution. Constraints on supportable use cases are guided by the consideration of the Test of Friction Risk when outside the endorsed remit of the Technology and Data functions. Model Governance combines these dimensions into a transparent lifecycle from initial model design through validation and deployment, ongoing operation, and final retirement.



Fig 3: The Tri-Pillar Governance Model: A Multi-Dimensional Framework for Autonomous ML Validation and Control in Financial AML Ecosystems

5.2. Security, Resilience, and Fraud Prevention

Trust in any system encompasses the belief that it will fulfill its intended purpose without compromise. In the context of cross-border finance, factors beyond model decisions contribute to risks faced when connecting with counterparties. Malicious actors may interfere with arrangements made by intelligent agents, affecting the outcomes of transactions or deliveries. Adversarial data may result in unacceptable model responses. A lack of data integrity during the operational use of generative models

may yield unacceptable products. Therefore, risk controls for cyber threats, adversarial model stealth, and data integrity when generating content are paramount. Incidents affecting the successful delivery of any element of a financial transaction require dedicated playbooks. Relevant actors must plan for interruptions in the facilitating technology, considering the operational needs of parties at all points of contact.


Cyber risk controls encompass the implementation of technical controls in the user, model provider, and model consumer environments to protect the confidentiality, integrity, and availability of the supporting infrastructure. Resilience to adversarial perturbations reducing the performance of a model beyond acceptable tolerances should be integrated into the architecture. Models should be required to pass adversarial robustness tests before deployment. Generated data should undergo appropriate validation and testing to confirm adherence to expected use. Any activity producing generative data, whether the input or content generated, is flagged for independent scrutiny both during and after the operational phase to prevent the use of such actors or products within any financial transaction. Transaction continuity playbooks for third-party model providers and support duties for synthetic asset swaps should also be agreed upon.

6. Generative AI and the Transformation of Financial Trade and Settlement

The emergence of autonomous generative AI alters the nature of trading and settlement in capital markets, encompassing market-making, pricing strategies, liquidity provision, and compliance considerations. These aspects transition from being the province of human traders to being autonomously governed by purpose-prepared agents. Market-making generates profits to the extent that the spreads earned on risk-neutral transactions are greater than losses on transactions with mispricing. Profits arising from static pricing and market-implied pricing, the latter feeding off arbitrage trading, add to pricing risk. Strategies for

liquidity provision and pricing in capital markets that exploit these opportunities have been identified. Under AI autonomy, the transparency of pricing strategies takes on a new dimension. Just as the use of natural-language processing for automated financial journalism leads to the publication of investment-relevant news items by multiple news outlets on important market-moving events, and potentially diminishes their individual impact, so though the proprietary influences of human traders may be suppressed, the underlying model-based pricing approaches remain fully visible. These considerations pose important questions regarding the adequacy of innovation for generating sustainable liquidity over timeframes of importance to market participants. Testing the adequacy of liquidity provision and pricing liquidity require a lens that incorporates balance sheet variables along with the appropriate market variables.

The execution of trades also comes under the domain of autonomous AI agents. Clearly, trading and settlement protocols can be designed that implement compliance checks for the transaction-exposing parties and require record-keeping for submitted orders that support future audit activities. With all trade confirmations preserved in a machine-readable format, intelligent agents that speak the "distinctly unnatural language" of the regulatory code are enabled to perform automated compliance checks and file automated compliance reports.



Equation C: Governance “Tri-Pillar” tests as measurable constraints

The explicitly defines a governance framework with **Test of Quality, Test of Control, Test of Intelligence**.

C1) Turn the three tests into a gating function

Let:

- $Q \in [0,1]$: quality test score (validation/performance)
- $C \in [0,1]$: control test score (auditability, monitoring, change control)
- $I \in [0,1]$: intelligence test score (robustness, capability boundaries)

A strict “go-live” rule:

$$\text{Deploy} = \mathbf{1}[Q \geq q_{min}] \cdot \mathbf{1}[C \geq c_{min}] \cdot \mathbf{1}[I \geq i_{min}]$$

6.1. Market-Macing, Pricing, and Liquidity Provision

Autonomous machine learning reshapes market-making, pricing strategies, and liquidity risks. Market-making involves constant quotes on both sides of an asset at acceptable levels for buying and selling volumes. Unlike



non-AI strategies, the provision of liquidity to both sides cannot be based solely on currency holding and risk capacity because of frequent price formations, rapid execution, and often no-routed settlements. With target optimization, intelligent liquidity flow estimation, and pre-trade price dispersion analysis, it becomes possible for ML solutions to sense the liquidity offered with the costs for other participants. Pricing becomes a risk-hedged allocation of the cost for identifying liquidity in real-time before profit utilization through non-routed instant transactions or through orders placed on more liquid platforms. Trading models can finally find their independent and complex auto-generated application in the properly managed ML and in the continuously adjusted market understanding. An original-systematic market impact analysis is determined in terms of buy-sell ratio or based on other ML criteria at different levels, where Market Impact is represented by the relation between the flow and the bought-sold currency at the execution spot. The decision is based on costs-effects-execution relations generated for the specific conditions of the instrument monitored by the strategy. In the absence of clear expected cost-benefit utility the trade is ruled out. The determination of the usable effect on liquid rate normality can largely improve the sustainability of the market-making process.

6.2. Compliance-Aware Trading and Settlement Protocols

By leveraging autonomous machine-learning, CI may autonomously notify the relevant markets of its intent to buy or sell a financial instrument or instrument basket, while specifying the price range and timing of the proximity of the intended buy/sell, with the goal of managing BIC's liquidity within the specified parameters. The proposed action is then visible to all participants in the same ecosystem. Before executing their buy/sell order, market autonomies that are thus alerted to such an intensity (and to the possible impact on the price of the instrument) could trigger an audit of "compliance" of the near-future action proposed by CI by CI or welcomed by CI. The audit

turns on whether the near-future action globally fulfills previous unsatisfied KYC/AML or other regulation alerts. Execution of the buy/sell order by CI would be allowed only if all the technical audit conditions are fulfilled. Compliance check could also be imposed on other market autonomies located in other ecosystems whose execution of the order would affect the price of the instrument and any selected nearby compliance criteria. Thus, during the buy/sell order process period, regulation alerts non-fulfilled and indicators show compliance issues not correctly coped with could also automatically trigger an audit on the execution of the buy/sell order under any compliance aspects.

Next, AI trading strategies—using reinforcement learning techniques—should be able not only to fulfil their trading mandate but also to be compatible and aligned with the trading community being offered or suggested—a key requirement for liquidity generation and efficient market-making. The alignment could concern any trading aspect (e.g. market impact, technological level, etc.). Such constraints could even anticipate executed trades for being already offered to the AI performing the trades, achieving a better price for the same market impact.

7. Conclusion

In conclusion, executing autonomous finance across borders requires governance frameworks that address the challenges presented by technology and markets. Supporting objectives, stakeholders, and principles have been identified, covering transparency, explainability, accountability, and alignment across jurisdictions. Surveillance and risk management requirements are tightening, focused on detection and prevention rather than enforcement, underpinned by classic KYC/AML-type data. A risk-based, outcome-oriented approach is emerging, driven by the need to balance innovation and public safety. As generative AI tools become more widely available, automatic, data- and model-consistent checks of regulatory compliance at the review-and-conformance stage of the trade life cycle are expected to provide a further layer of

confidence and deterrence.




Fig 4: Pillars of Cross-Border Governance

The responsible use of generative AI is expected to further enhance efficiency but concentrated liquidity poses greater risk. Under conditions of normality or uncertainty, such an imbalance may not threaten stability. A lapse of data or model quality or an accumulation of adversarial attacks have the potential of creating a gaping hole that traders and market-makers are unable to pre-emptively fill. While the resultant pricing may still pass the gut test of several market experts—those traders and market-makers operating on discretion—whether such pricing is value-relevant is a different matter. Two considerations remain pertinent for the future of autonomous finance across borders, particularly within Europe. First, the debate around an appropriate regulatory framework governing financial trade systems based on generative AI is still at a relatively nascent stage. Second, the challenge of appealing to a variety of AI-augmented techniques that could make possible the desired execution of a conforming protocol on public blockchain technology remains to be fully addressed.

7.1. Final Insights and Future Directions in Autonomous Finance

The insights support a risk-based and outcome-oriented approach to advance autonomous finance across borders, laying the foundation for a regulatory and supervisory ecosystem that addresses governance, compliance, and the increasing role of generative AI in cross-border financial interactions. Sovereignty over data and the integration of international standards with national regulation are paramount. Friction resulting from divergence across jurisdictions can be reduced through harmonization of requirements, mutual recognition agreements, cross-border data-sharing arrangements, and greater supervisory coordination. Resourceful and resilient organizations will be the first choice of the regulatory community. It remains crucial to keep issues of outcome and impact in sight, regardless of the level of autonomy attained. Generative AI is expected to play an increasingly prominent role in the autonomous processing of trades and settlement over time, driving improvements in areas such as market-making pricing and liquidity strategies and compliance-aware trading protocols. However, lack of control and monitoring continues to be a concern, not least when the underlying models remain opaque. Consequently, compliance, risk management, and control have yet to become fully embedded in generative AI-driven financial services, necessitating the development of a framework that maps regulatory expectations to AI-enabled processes. Human intervention will eventually be required to ensure appropriate context, prudence, and soundness for high-risk actions in a risk-based model-as-a-safety-net approach to governance of both advisory and autonomous actions.

8. References

- Alessi, L., & Savona, R. (2021). Machine learning for financial stability. In S. Consoli, D. Reforgiato Recupero, & M. Saisana (Eds.), *Data science for economics and finance* (pp. 65–87). Springer.
- Liu, L., Chen, C., & Wang, B. (2022). Predicting financial crises with machine learning methods. *Journal of Forecasting*, 41(8), 1684–1701.



3. Danielsson, J., Macrae, R., & Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance*, 140, 106290.
4. Yandamuri, U. S. (2024). AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology. *Metallurgical and Materials Engineering*.
5. Mirestean, A. T., & Ravikumar, R. (2021). Powering the digital economy: Opportunities and risks of artificial intelligence in finance. *International Monetary Fund*.
6. OECD. (2021). Artificial intelligence, machine learning and big data in finance: Opportunities, challenges and implications for policy makers. OECD Publishing.
7. OECD. (2021). OECD business and finance outlook 2021: AI in business and finance. OECD Publishing.
8. Schmitt, L. (2022). Mapping global AI governance: A nascent regime in a fragmented landscape. *AI and Ethics*, 2, 303–314.
9. Gill, A. S., & Germann, S. (2022). Conceptual and normative approaches to AI governance for a global digital ecosystem supportive of the UN Sustainable Development Goals (SDGs). *AI and Ethics*, 2, 293–301.
10. Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. *AI and Ethics*, 2, 603–609.
11. Wirtz, B. W., Weyerer, J. C., & Kehl, I. (2022). Governance of artificial intelligence: A risk and guideline-based integrative framework. *Government Information Quarterly*, 39(4), 101685.
12. Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2022). Toward AI governance: Identifying best practices and potential barriers and outcomes. *Information Systems Frontiers*, 25, 123–141.
13. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. *Journal of Neonatal Surgery*, 13(1), 2387-2399.
14. Mökander, J., & Floridi, L. (2023). Operationalising AI governance through ethics-based auditing: An industry case study. *AI and Ethics*, 3, 451–468.
15. Černevičienė, J., & Kabašinskas, A. (2022). Review of multi-criteria decision-making methods in finance using explainable artificial intelligence. *Frontiers in Artificial Intelligence*, 5, 827584.
16. Doumpos, M., Zopounidis, C., Gounopoulos, D., Platanakis, E., & Zhang, W. (2023). Operational research and artificial intelligence methods in banking. *European Journal of Operational Research*, 306(1), 1–16.
17. Ahmed, I. E., Mehdi, R., & Mohamed, E. A. (2023). The role of artificial intelligence in developing a banking risk index: An application of adaptive neural network-based fuzzy inference system (ANFIS). *Artificial Intelligence Review*, 56, 13873–13895.
18. Kalyani, S., & Gupta, N. (2023). Is artificial intelligence and machine learning changing the ways of banking: A systematic literature review and meta analysis. *Discover Artificial Intelligence*, 3, 41.
19. Mangalampalli, B. M. (2024). AI-Enhanced Data Governance: Automating Compliance In Healthcare Analytics Platforms. *The Review Of DIABETIC STUDIES OPEN ACCESS*.
20. Weber, P., Carl, K. V., & Hinz, O. (2024). Applications of explainable artificial intelligence in finance—A systematic review of finance, information systems, and computer science literature. *Management Review Quarterly*, 74, 867–907.
21. Chen, X.-Q., Ma, C.-Q., Ren, Y.-S., Lei, Y.-T., Huynh, N. Q. A., & Narayan, S. (2023). Explainable artificial intelligence in finance: A bibliometric review. *Finance Research Letters*, 56, 104145.
22. Giudici, P., & Raffinetti, E. (2023). SAFE artificial intelligence in finance. *Finance Research Letters*, 56, 104088.
23. Kovvuri, V. R. R., Fu, H., Fan, X., & Seisenberger, M. (2023). Fund performance evaluation with explainable artificial intelligence. *Finance Research Letters*, 58, 104419.



24. Zhou, Y., Li, H., Xiao, Z., & Qiu, J. (2023). A user-centered explainable artificial intelligence approach for financial fraud detection. *Finance Research Letters*, 58, 104309.
25. Rodgers, W., Hudson, R., & Economou, F. (2023). Modelling credit and investment decisions based on AI algorithmic behavioral pathways. *Technological Forecasting and Social Change*, 191, 122471.
26. Pamuk, M., Schumann, M., & Nickerson, R. C. (2024). What do the regulators mean? A taxonomy of regulatory principles for the use of AI in financial services. *Machine Learning and Knowledge Extraction*, 6(1), 143–155.
27. Pan, M., Li, D., Wu, H., & Lei, P. (2024). Technological revolution and regulatory innovation: How governmental artificial intelligence adoption matters for financial regulation intensity. *International Review of Financial Analysis*, 96, 103535.
28. Azzutti, A. (2024). AI governance and algorithmic trading: Some regulatory insights from the EU AI Act. *Banking and Finance Law Review*, 39, 133–168.
29. Crisanto, J. C., Leuterio, C. B., Prenio, J., & Yong, J. (2024). Regulating AI in the financial sector: Recent developments and main challenges. *Bank for International Settlements, Financial Stability Institute*.
30. Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., & Stein, M. (2024). Intelligent financial system: How AI is transforming finance. *Bank for International Settlements, BIS Working Papers No. 1194*.
31. Leitner, G., Singh, J., van der Kraaij, A., & Zsámboki, B. (2024). The rise of artificial intelligence: Benefits and risks for financial stability. *European Central Bank, Financial Stability Review*, 1, 1–13.
32. Financial Stability Board. (2024). The financial stability implications of artificial intelligence. *Financial Stability Board*.
33. Mangala, N. (2024). Leveraging Microsoft Fabric lakehouse as an AI-ready data platform for enterprise analytics. *Journal of Information Systems Engineering and Management*.
34. OECD. (2024). Regulatory approaches to artificial intelligence in finance. *OECD Artificial Intelligence Papers*, No. 24. OECD Publishing.
35. Bonaparte, Y. (2024). Artificial intelligence in finance: Valuations and opportunities. *Finance Research Letters*, 60, 104851.
36. Beckmann, L., & Hark, P. F. (2024). ChatGPT and the banking business: Insights from the US stock market on potential implications for banks. *Finance Research Letters*, 63, 105237.
37. Kirtac, K., & Germano, G. (2024). Sentiment trading with large language models. *Finance Research Letters*, 62, 105227.
38. Almeida, J., & Gonçalves, T. C. (2024). The AI revolution: Are crypto markets more efficient after ChatGPT 3? *Finance Research Letters*, 66, 105608.
39. Klein, T., & Walther, T. (2024). Advances in explainable artificial intelligence (xAI) in finance. *Finance Research Letters*, 70, 106358.