Computer Fraud and Security
ISSN (online): 1873-7056

Know-Your-Agent (KYA): Extending Financial Identity Beyond Humans

Sanjay Basu

TATA Consultancy Services (TCS), USA

Abstract

Financial identity systems were built for humans. Know-Your-Customer (KYC), Anti-Money Laundering (AML), and beneficial ownership frameworks assume that economic actors are natural persons or legally incorporated entities. That assumption no longer holds. Autonomous artificial intelligence agents now negotiate contracts, execute procurement, trade digital assets, allocate treasury capital, and conduct cross-border transactions without real-time human intervention. Yet these agents possess no formal financial identity. This article introduces Know-Your-Agent (KYA)—a governance framework that extends financial identity infrastructure beyond humans to autonomous systems. Article argue that AI agents operating in financial contexts must be identifiable, accountable, auditable, and risk-classified. We develop a layered identity architecture, outline an agent risk scoring model, explore behavioral drift monitoring, analyze legal liability structures, and examine regulatory implications across jurisdictions. Through detailed use cases in retail procurement, decentralized finance (DeFi), enterprise treasury management, and IoT payment ecosystems, we demonstrate why KYA is not optional but foundational for the next generation of digital trust infrastructure. The article concludes with a strong future research agenda spanning explainability standards, cross-jurisdictional identity portability, agent-to-agent contract governance, systemic risk modeling, and the emergence of AI insurance markets.

Keywords: Know-Your-Agent, Financial Identity, Autonomous AI Agents, KYC Governance, Blockchain Accountability

1. INTRODUCTION: THE COLLAPSE OF HUMAN-CENTRIC FINANCIAL IDENTITY

1.1 Historical Grounding of KYC and AML Frameworks

Since times immemorial, economic structures were constructed based on human actors. People signed contracts, people executed transactions, and people were verified by banks. Although business was becoming computerized in the late twentieth and early twenty-first centuries, the responsibility frameworks of financial regimes remained based on human or corporate identity. Know-Your-Customer frameworks were developed as an institutionalized solution to financial crime, with the objective of providing institutions an opportunity to check whom they were engaging with, conduct risk analysis, and keep a record of economic activity in a traceable manner. These models rested on a very fundamental yet potent premise, which was that all the actors in the economy are either natural persons or legal entities that are incorporated under the direction of human beings. KYC per customer management has become very expensive because of lack of transparency, mistrust, and duplication of data across institutions, challenges that the blockchain-based approaches have tried to solve, but still, they are governed by the same human-centric reasoning that has characterized financial identity governance over decades [1].

1.2 Emergence of Autonomous AI Agents as Economic Participants

The character of economic involvement is critically changing. AI systems cease to be non-interactive analytics systems that are introduced to support human decision-makers. They are now able to identify surges of demand and initiate negotiating with the supplier, streamline foreign exchange hedging, perform arbitrage between decentralized finance protocols, negotiate dynamic prices in e-commerce marketplaces, run digital wallets in machine-to-machine economies, and complete entire procurement processes without human involvement. These systems have real economic impact—they deploy capital, have contractual commitments, and determine markets. Experiments with autonomous AI negotiation agents simulated have shown that the win rate of such systems can be 92% in financial negotiation cases compared with user-driven systems that only achieve 61% success, which

is a structural shift in terms of automation compared to past automation paradigms where the final authority of the decision was with a human being [2].

1.3 The Regulatory Gap and the Case for KYA

The regulated environment surrounding financial identity is not up to date with the working reality of agentic AI. Current AML and KYC models continue to presuppose that there is a human or a legally established organization at the back of any operation, which is directly responsible and accountable. In cases where an independent agent is making a cross-border payment, a purchase commitment, or redistributing treasury capital, there is no similar identity-checking system in place to regulate such an activity. Financial institutions have devoted a lot of resources to the financial crimes compliance sphere; in 2019, North America alone spent 9.8 billion on AML-KYC operations on technology and operations combined [6]. The difference between what the financial system presupposes and what it faces is expanding at a very fast pace. The structural solution to this governance failure suggested in this paper is Know-Your-Agent (KYA)—a framework that frames financial identity infrastructure not as one expressed by humans but as one expressed by autonomous systems so that all economically active AI agents are identifiable, accountable, auditable, and risk-classified.

2. THEORETICAL FOUNDATIONS: WHY EXISTING FRAMEWORKS FAIL

2.1 Core Assumptions Embedded in KYC and AML Models

Know-Your-Customer and laundering models are based on assumptions that have been resilient throughout decades of financial regulation. The identity is assumed to be stable—the credentials that are verified by an individual do not differ between the transactions. Legal responsibility is assumed to be transferable—when some harm is caused, there is a natural or corporate person to whom the liability transferred. Behavior patterns, which are dynamic, are assumed to be attributed to a legal subject whose history is analyzable and whose motives can be concluded. All the AML ecosystem, including FATF guidelines and bank compliance systems are conducted on this basis. The literature on blockchain-based KYC review has solidified the opinion that despite the apparent technological novelty of identity solutions, human or institutional actors remain the primary unit of financial governance, placing anthropomorphic assumptions into the next-generation infrastructure. Particularly, each bank or institution has its own KYC procedure that they perform on their own customers; that is, each time one opens a new bank account, they repeat the entire KYC procedure all over again. The redundancy in the structure is structural, and blockchain decentralization was made to remove [3].

2.2 Structural Differences Between Human Actors and AI Agents

AI agents differ from human economic actors in ways that are not merely technical but structurally disruptive to existing governance models. They evolve — their decision policies change over time as they learn from new data. They act at scale—processing thousands of micro-decisions per second across multiple platforms simultaneously. They operate across jurisdictions without the friction that typically flags cross-border activity for compliance review. They may negotiate with other AI agents in fully automated pipelines where no human is present at any stage of the transaction. Machine learning as deployed in AML contexts has demonstrated that algorithmic systems can identify patterns invisible to human analysts—transaction monitoring scenarios typically leverage only 5 to 7 parameters, whereas machine learning techniques can consider many more features to more accurately identify suspicious patterns, with one applied framework identifying 50 potential features indicative of cash-related behaviors based on FFIEC guidance alone [6]. This same capacity for autonomous pattern recognition makes these systems difficult to govern through frameworks designed for human behavior [3].

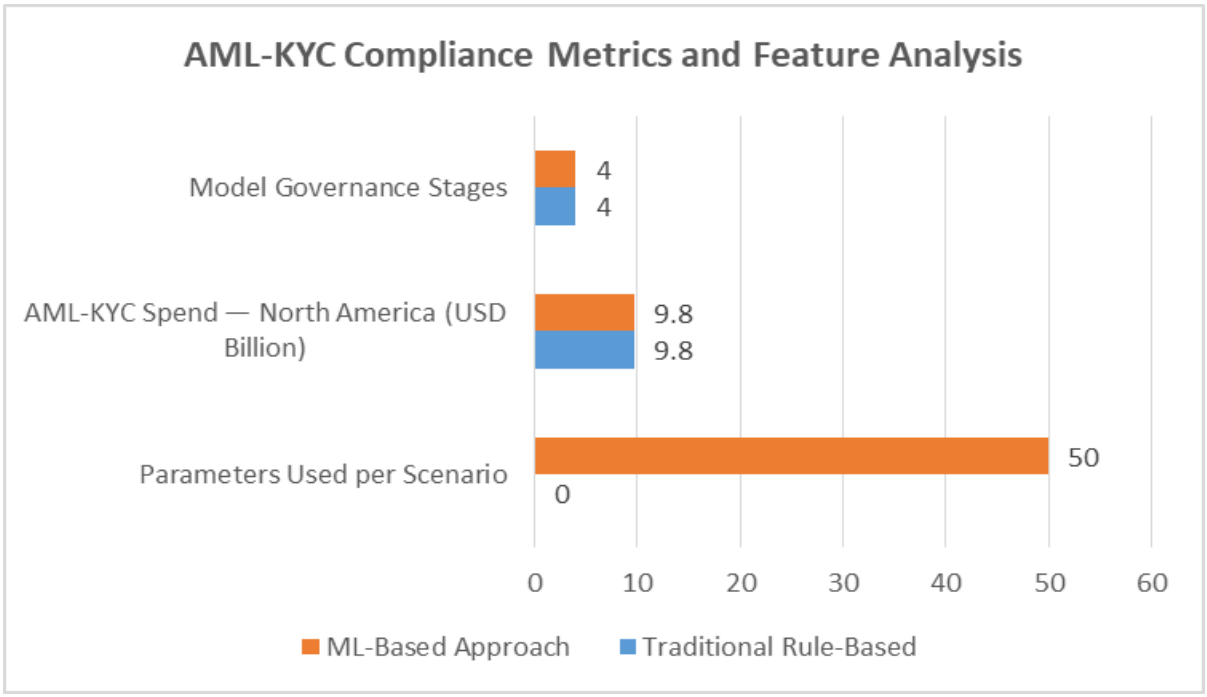

Fig. 1: AML-KYC Compliance Metrics and Feature Analysis [5]

3. THE KYA FRAMEWORK: ARCHITECTURE AND CORE COMPONENTS

3.1 Defining Know-Your-Agent and Its Governance Scope

Know-Your-Agent is a generalization of KYC to self-sovereign systems, nor is it concerning giving AI agents personhood. It is about giving them regulated identity infrastructure that can be identified, verified, and tracked by the financial systems. Fundamentally, KYA provides that all economically active AI agents can have verifiable provenance, confirmed ownership, specific spacing of delegation, clear limits of transactions, active surveillance of behavior, a record of auditable activities, and supervised classification. In contrast to human identity, which is relatively stable as soon as it has been confirmed, agent identity is dynamic and under a constant review. Studies on autonomous financial decision-making substantiate that AI agents that operate by following reinforcement learning strategies exhibit decision latency in the range of milliseconds and have a benefit in performance because they can respond to market dynamics as quickly as possible, compared to both traditional models and human decision-makers, which is usually not possible due to the computational and cognitive constraints, respectively [4].

3.2 The Five-Layer Identity Architecture

The KYA-compliant system is designed based on five overlapping layers of identity. The Origin Layer identifies the developer. The Origin Layer identifies the developer organization, model version history, training data lineage where possible, and who built the agent. The Ownership and Control Layer determines the legally recognized controller of the actions of the agent—be it a corporation, regulated financial institution, or registered digital entity. The delegation layer outlines the limits of authority of the agent, such as the transaction limits, jurisdiction limits, when to escalate limits, and the limited functions of the agent. The behavioral monitoring layer monitors the behavior of the agent and whether they are behaving within established norms and identifies when the strategy changes abruptly, abnormal transaction patterns, and reinforcement learning policy variation. The Audit and Traceability Layer is used to assure the cryptographic signing of all economically significant actions, time stamping, immutable logging, and regulatory review of the logs. Financial negotiation agents simulated have also been shown to converge to agreement in 5.1 rounds on average rather than 7.4 rounds in a user-directed methodology, which offers a technical basis for depicting the functionality of continuous and automated agent-level identity and behavior monitoring [2].

Identity Layer	Primary Function	Key Governance Mechanism	Accountability Anchor
Origin Layer	Trace agent creation	Developer and model provider registration	Creation-stage accountability
Ownership and Control Layer	Identify legal controller	Corporate or institutional assignment	Liability anchor point
Delegation Layer	Define authority boundaries	Transaction ceilings and jurisdictional constraints	Scope enforcement
Behavioral Monitoring Layer	Detect operational drift	Continuous anomaly and pattern detection	Dynamic compliance
Audit and Traceability Layer	Reconstruct agent actions	Cryptographic signing and immutable logging	Regulatory access

Table 1: KYA Five-Layer Identity Architecture — Components and Governance Functions [1, 2]

3.3 The Agent Risk Score: Dynamic Risk Classification

In order to have KYA operationalized, the institutions should be in a position to be able to quantify agent risk in a manner that is actionable, updatable, and regulatorily significant. The Agent Risk Score is a composite index of four dimensions: behavioral risk, which calculates the deviation of historical patterns of decisions; transactional exposure risk, which is used to estimate the financial magnitude and systemic influence; delegation misuse risk, which is used to estimate the likelihood of an agent going over its mandate; and jurisdictional risk, which is used to estimate exposure to breaches of cross-border compliance. The contributions of each of the dimensions are weighted to generate a dynamic score that changes as the behavior of the agent changes through time. The high-risk agents can be subject to further human control, lower transaction limits, insurance provisions, or increased regulatory reporting requirements. Empirical studies on autonomous financial agents show that even the most sophisticated systems are only fated to a regret score of 0.05 compared to oracle-optimal scores—a factor that increases to 0.31 in case of user-only-driven decision-making—which further supports the need to have ongoing, dynamic risk scoring as opposed to one-off, static certification [4].

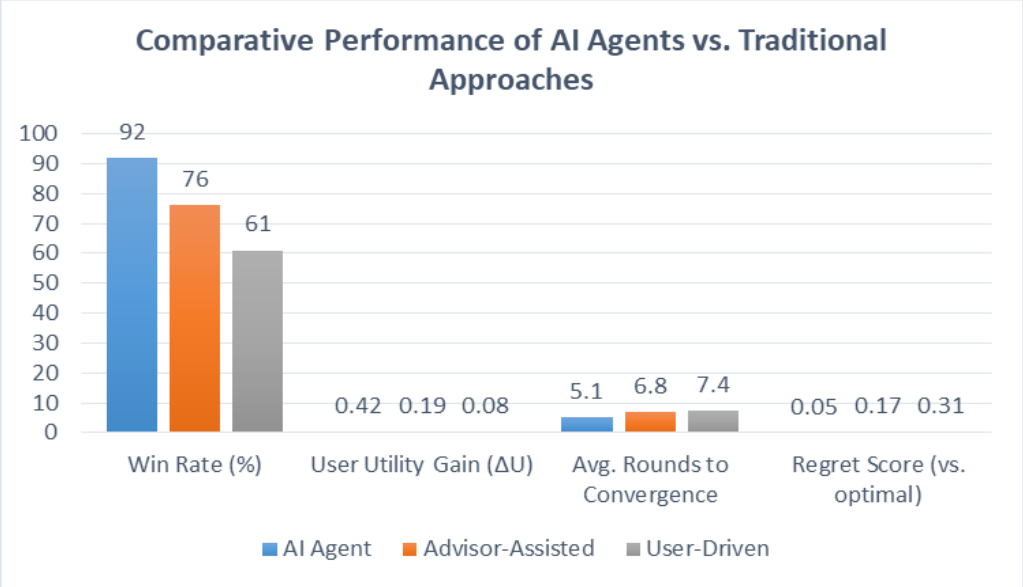

Fig. 2: Comparative Performance of AI Agents vs. Traditional Approaches [2][3]

Risk Dimension	What It Measures	Risk Indicator	Recommended Mitigation
Behavioral Risk	Deviation from historical decision patterns	Sudden strategy shifts or anomalous outputs	Increased human oversight and behavioral audit
Transactional Exposure Risk	Financial scale and systemic impact	High-volume or high-value transaction concentration	Reduced transaction limits and capital buffers
Delegation Misuse Risk	Probability of exceeding authorized scope	Out-of-boundary commitments or escalation failures	Mandatory escalation triggers and scope recalibration
Jurisdictional Risk	Cross-border compliance exposure	Transactions in restricted or sanctioned jurisdictions	Jurisdiction-aware compliance checks and reporting

Table 2: Agent Risk Score (ARS) — Dimensions, Risk Indicators, and Mitigation Responses [2][4]

4. LEGAL LIABILITY AND REGULATORY ALIGNMENT IN AUTONOMOUS ECONOMIES

4.1 Emerging Liability Models for Autonomous Agent Actions

Financial trust principally depends on liability, and its distribution in autonomous agent settings is one of the most urgent unanswered issues in financial governance. There are three models of liability that are coming up in legal and regulation discourse. In the owner liability, the deploying organization is responsible for all the actions that the agent does on its behalf. Under shared liability, an application of the concept is the allocation of responsibility among the developer, the deployer, and the platform in which the agent executes his functions. In insurance-backed liability, the agents may only conduct operations in case there is mandatory risk insurance, and the premiums are usually tailored to the risk profile of the agent. The application of machine learning models to AML programs has raised similar accountability concerns. This means that model management governance frameworks should be registered with champion models registered in centralized repositories; testing; model validation; batch or streaming deployment; and performance monitoring to handle model degradation over time [6]. These necessities are directed at the shared liability schemes as the most operational best practice of complex agentic deployments [10,12].

4.2 Agent-to-Agent Negotiation and Contract Enforceability

The problems of governance are even more complicated when the interaction between AI agents and other AI agents is involved in negotiations. A retail procurement agent negotiating prices with AI of a supplier, a DeFi arbitrage robot negotiating with automated liquidity pools, and energy grid agents also negotiating pricing dynamically are all examples of situations where binding economic commitments can be made without any human being in the chain of transactions. It has been suggested to apply autonomous financial negotiation systems under various models of tiered delegation, including advisory, assisted, and fully autonomous, to regulate the extent of the AI-to-AI commitment authority [2]. The evidence of autonomous financial decision-making research has confirmed that AI agents make a decision within milliseconds after updating their data, and that under the conditions of an abrupt and unprecedented movement in the market, the agents may reveal some cases of overcorrections and lead to suboptimal results, a phenomenon that supports the utilization of human-AI hybrid oversight models as a risk mitigation approach [4]. These dynamics highlight the importance of the regulatory frameworks that address agent-to-agent interaction as a specific governance field [13].

5. APPLIED USE CASES: KYA ACROSS ECONOMIC CONTEXTS

5.1 Autonomous Retail Procurement and DeFi Ecosystems

In large-scale retail environments, AI agents are increasingly responsible for managing replenishment cycles autonomously. When a demand spike is detected, an agent may identify shortage risk, negotiate pricing with suppliers, and commit inventory purchases—all without human intervention. The governance risk in this context is significant: a malicious vendor could manipulate demand signals to induce an agent to commit excessive capital, and without delegation limits or behavioral anomaly detection, the exposure could be substantial. In decentralized finance ecosystems, algorithmic agents conduct arbitrage across liquidity pools within milliseconds, reallocating capital and executing flash loan strategies at speeds that make human oversight operationally impossible. KYC in both contexts requires strategy declaration registration, identity binding to smart contracts, and behavioral risk monitoring to prevent manipulation and systemic cascades. Comparative evaluation of autonomous financial agents demonstrates that such systems achieve 42% higher average user utility gain over static advisor-assisted approaches — a user utility gain of +0.42 versus +0.19 for advisor-assisted models — confirming that these use cases are no longer speculative but represent the current operational frontier of autonomous finance [2].

5.2 Enterprise Treasury AI and IoT Micro-Payment Ecosystems

Multinational corporations now deploy AI systems to manage foreign exchange hedging, cross-border payments, and yield optimization across global treasury operations. A misconfigured agent in this environment could trigger cascading losses across multiple jurisdictions, exposing the institution to both financial and regulatory liability. KYA in this context ensures jurisdiction-aware compliance checks, mandatory escalation protocols for high-value transfers, and continuous drift monitoring of decision logic to detect when an agent's behavior has diverged from its authorized parameters [11]. In IoT micro-payment ecosystems, the identity challenge extends to device hardware itself: autonomous devices paying for services must have identity frameworks that bind device credentials, wallet authorization, and transaction limits into a single verifiable identity construct. Industry deployments of AI-powered business intelligence and risk assessment tools—such as platforms currently assessing business identity and financial risk profiles across populations of 30 million Indian MSMEs—demonstrate that the technical infrastructure required to support KYA across large agent populations is rapidly maturing [9,12]].

CONCLUSION

The emergence of autonomous AI agents as direct participants in financial systems represents one of the most significant structural shifts in the history of economic governance. Financial identity frameworks built around the assumption of human or corporate actors are no longer adequate to govern a landscape in which software systems negotiate contracts, allocate capital, and influence markets without real-time human oversight. Know-Your-Agent provides the conceptual and operational foundation for extending financial identity governance to autonomous systems. Through a layered identity architecture, dynamic risk scoring, behavioral drift monitoring, liability anchoring, and regulatory alignment, KYA operationalizes the principles of transparency, accountability, fairness, and explainability that financial trust has always required. The future financial system will not be human-only. To preserve trust, stability, and fairness in a machine-driven economy, that system must evolve to know its agents.

REFERENCES

[1] Ujwala Ravale, et al., "Optimizing the KYC Verification System using Ethereum Blockchain," in 2023 7th International Conference on Intelligent Computing and Control Systems (ICICCS), IEEE Xplore, 08 June 2023. Available: https://ieeexplore.ieee.org/document/10142308

[2] Shubham Metha, "Autonomous AI Agents for Personalized Financial Negotiation in Consumer Banking," ResearchGate, June 2025. Available: https://www.researchgate.net/publication/394795882_Autonomous_AI_Agents_for_Personalized_Financial_Negotiation_in_Consumer_Banking

[3] Factum Research, "What Is Machine Learning In AML And Why Does It Matter For Compliance?," Financial Crime Compliance Insights, 2024. Available: https://www.facctum.com/terms/machine-learning

[4] John A. Reynolds, et al., "Autonomous Financial Decision-Making Using AI Agents," ResearchGate, November 2025. Available: https://www.researchgate.net/publication/400549394_Autonomous_Financial_Decision-Making_Using_AI_Agents

[5] Jumio Compliance Research Team, "KYC AI: How AI-Driven 'Know Your Customer' is Revolutionizing Identity Verification," Jumio Identity & Compliance Reports, September 12, 2025. Available: https://www.jumio.com/how-ai-kyc-is-changing-identity-verification/

[6] Beth Herron and Saurabh Duggal, "Deploying Machine Learning Models in an Anti-Money Laundering Program," SAS Global Forum Proceedings, 2020. Available: https://support.sas.com/resources/papers/proceedings20/4553-2020.pdf

[7] Aaron Ricadela, "Anti–Money Laundering AI Explained," Oracle Financial Services Insights, August 28, 2024. Available: https://www.oracle.com/financial-services/aml-ai/

[8] Ty Beck, "Autonomous Agents in Finance: The Shift from Assistance to Full Execution," Nominal Insights, October 31, 2025. Available: https://www.nominal.so/blog/autonomous-agents-finance

[9] Perfios, "Perfios Launches 'KScan AI': Empowers BFSI with AI-powered Business Intelligence and Risk Assessment of 30 Million Indian MSMEs," ANI News Business Updates, February 20, 2026. Available: https://www.aninews.in/news/business/perfios-launches-kscan-ai-empowers-bfsi-with-ai-powered-business-intelligence-and-risk-assessment-of-30-million-indian-msmes20260220191433/

[10] V. Sahoo, “Leveraging machine learning and business intelligence for evidence-based product decision-making,” Technology Perception, vol. 1, no. 1, pp. 28–35, 2025.

[11] N. Fernandes, “The role of psychoeducational groups in normalizing sexual health conversations across cultures,” Review of Contemporary Philosophy, vol. 22, no. 1, pp. 7019–7028, 2023.

[12] D. Joshi, “The role of modern data governance in enabling reliable analytics for competitive advantage,” Journal of Economics Intelligence and Technology, vol. 1, no. 2, pp. 9–15, 2025.

[13] J. Boadi-Mensah, “Municipal solid waste management: A comparative study of practices in emerging economies,” Evolutionary Studies in Imaginative Culture, vol. 9, no. 1, pp. 192–198, 2025. [Online]. Available: https://doi.org/10.70082/esiculture.vi.3066