# Breakthrough Autonomous Agentic AI Frameworks for Real-Time Multi-Counterparty Derivatives Orchestration: Self-Adaptive Multi-Agent Coordination for Enterprise-Scale Trading and Collateral Management

1st Appa Rao Nagubandi

*Lead Software Engineer,*ORCID ID: 0009-0005-8424-7071

***Abstract*—**Toward a Breakthrough Autonomous Agentic Framework A general-purpose and au- tonomous agentic framework offers breakthrough ca- pabilities for real-time multi-counterparty derivatives trading and trading-related derivatives collateral or- chestration and management. Its practical usefulness derives from self-adaptive coordination mechanisms based on multiple online learning cycles. Such frame- works are necessary for schelling's equilibrium choice problem in real-time derivatives trading and the in- telligent orchestration of systems in which trading is not just structurally multi-party but also in the log- ical sense and involves evolving orders that generate demand for additional systems. Such functions go well beyond those of a facilitator or a bank. Their inher- ent significance extends to systematic trading-related derivatives collateral management since adoption in- volves the participation of multiple counterparties. The general objectives are to design and evaluate empiri- cal trading-related schemes that showcase multi-party derivatives Orchestration capabilities and multi-agent systems from a practical viewpoint. It is assumed that each derivative Position has been entered in an exist- ing real-time neutral trading framework, ladies receive products with fair default risk and can be transacted in bulk, and all practical considerations important to the counterparts and the infrastructure operator have been sufficiently stabilized. Within this narrow scope, the focus is on elaborating achievable general-purpose order-execution strategies with guarantees of collateral demand and economic coverage.

***Index Terms*—**Autonomous Agentic Framework, MultiCounterparty Trading, Derivatives Orchestra- tion, RealTime Trading Systems, Online Learning Cycles, SelfAdaptive Coordination, Schellings Equi- librium Modeling, Intelligent System Orchestration, MultiAgent Trading Models, Trading Infrastructure Design, Derivatives Collateral Management, Counter- party Risk Handling, Neutral Trading Frameworks, Bulk Derivative Transactions, OrderExecution Strate- gies, Economic Coverage Guarantees, Collateral De- mand Optimization, RiskNeutral Trade Design, Empir- ical Trading Schemes, RealTime Market Intelligence.

Identify applicable funding agency here. If none, delete this.

## INTRODUCTION

Algorithmic solutions power increasingly complex trad- ing operations. The lack of coordination in real-time trading between multiple bid and ask quotes (by mul- tiple counterparties), nevertheless, poses a challenge to the emergence of autonomous trader agents that fully complete the logical technology stack without the need for human supervision. Such capabilities would provide risk-mitigated self-clouding across multiple counterparties, including venues, at an enterprise scale. Contracts engaged directly with liquidity providers, which directly assume the corresponding risk, should allow agencies to fulfil their mission. Breakthrough self-adaptive, self-correcting, multi-agent enterprise agentic frameworks for independent traders could, therefore, transform derivatives and other transactions involving multiple counter- parties into a fully automated process capable of real-time collaborative execution without centralised control. A logical framework encompassing the full technology stack (including enter- prise agent capabilities) should facilitate direct engage- ment with liquidity-providing counterparties and allow the fully autonomous completion of indirectly engaged orders. The capacity to fulfil self-clouding independently would thus mitigate risk at network level. Moreover, enabling independent margin-matched transaction cloud- ing would allow venue-agnostic real-time trading in any derivative, including exotic contracts, via bids with visibly multiple layers of execution risk clearly visible in the limit order book. Mutual benefit and protective guarantees for non- risk-taking venues are thus realised. Socialising risk on hedges should also permit direct engagement with venues that normally hedge order execution via referential pricing alone.

### A. Overview of the Autonomous Agentic Framework

A breakthrough autonomous agentic framework enables real-time multi-counterparty orchestration of the full life-cycle of derivatives trades and collateral management via


Fig. 1. Autonomous Agents: Real-Time Derivatives & Collateral Orchestration

self-adaptive multi-agent coordination. The cross-lifecycle orchestration of derivatives trades and the management of collateral held in connection with these trades across multiple counterparties remain challenging problems. The work addresses aspects of both problems by providing a framework that enables how a financial institution such as a bank, pension fund, or insurance company can adaptively orchestrate these operational functions across multiple counterparties in real time using decision agents, execution agents, risk agents, compliance agents, and rec- onciliation agents. A key part of the novelty lies in the provision of self-adaptation mechanisms that enable the specification and adaptation of preferences, constraints, and decision policies in response to changes in the trading environment and impact-evaluation feedback loops that support stability. The research satisfies an acknowledged need for a practical framework that supports real-time multi-counterparty derivatives orchestration. Demonstrat- ing this capability and hence its practical relevance marks a major milestone towards enabling the development and deployment of self-adaptive autonomous agentic systems for the real-time trading of derivatives and associated management of collateral that is necessary for enterprise- scale operations.

their orders for trade and the necessary collateral guarantees (covering margin requirements for the trade itself and for other counterparty risk) in a way that is self-adaptive and responsive to the trading environment. The need for the proposed autonomous agentic framework to close the latency gap in derivatives trading remains a critical element behind many actors (exchange venues, ATSs, brokers, and banks). Consequently, the related properties must be further articulated and captured in the implementation. The area of real-time derivatives trading across multiple counterparties involves several major actors across the financial industry ecosystem, all engaging in the respective business processes under strict requirements for performance, security, compliance (legal and anti-money laundering) and risk management. An investment house would need to cover all of the aspects of the process when trading into an instrument expiring much later in time than into the execution point, where the bank's borrow rate and collateral consequences would be of concern but also regulatory criteria, such as those defined for clearinghouses, for fundamental liquidity; and, if the ARCH theory holds, concealed liquidations in case of strong downturns. These actors would need to trade into the instrument at a wide enough price difference to satisfy all of these costs, since the exchanges are required to be honest facilitators of price discovery and the pitch's pricing environment correctly compute the liquidity, concentration, funding and horizon limits above, after accounting for the margin and other collateral requirements on the anticipated trades.

**Equation 1 – Multi-Agent Joint Policy Optimiza- tion**

**Deriving the joint policy gradient**

We want $\nabla_{\theta_i} J(\theta)$ for each agent.

1. Write the trajectory probability:

$$p_\theta(\tau) = \rho(s_0) \prod_{t=0}^{\infty} \pi_\theta(a_t \mid s_t) P(s_{t+1} \mid s_t, a_t) \quad (1)$$

2. Rewrite $\nabla_\theta J(\theta)$ as expectation over trajectories:

$$J(\theta) = \int p_\theta(\tau) R(\tau) d\tau \quad (2)$$

where $J(\theta) = \sum_{t=0}^{\infty} \gamma^t r_t$

I. PROBLEM DOMAIN AND REQUIREMENTS

The functional and nonfunctional requirements for

1. Differentiate under the integral (likelihood-ratio trick):

$$\nabla J(\theta) = \int \nabla p(\tau) R(\tau) d\tau = \int p(\tau) \nabla \log p(\tau) R(\tau) d\tau$$

autonomous agentic systems in real-time derivatives $\theta_i$ trading across multiple counterparties. Real-time derivatives trading across multiple counterparties

$\theta_i$ $\theta$

$\theta$ $\theta_i \theta$(3)

requires functional and nonfunctional requirements to be satisfied, including the ability to orchestrate a set of agents that autonomously take control of decision-making and execution processes (in a limited sense), negotiating, coordinating, and executing with other counterparties

2. Compute $\log p_\theta(\tau)$. Only the policies depend on $\theta$:

$$\log p_\theta(\tau) = \log \rho(s_0) + \sum_{t=0}^{\infty} [\log \pi_\theta(a_t \mid s_t) + \log P(s_{t+1} \mid s_t, a_t)] \quad (4)$$

| time | exposure | margin_posted | net_risk |
|---|---|---|---|
| 0 | 0.0 | 0.0 | 0.0 |
| 1 | 1.764052345967664 | 0.882026172983832 | 0.882026172983832 |
| 2 | 1.8113990851413546 | 0.9056995425706773 | 0.9056995425706773 |
| 3 | 2.4278572 | 1.21392862610 | 1.21392862 |

| | | | |
|---|---|---|---|
| | 52218823 | 94115 | 61094115 |
| 4 | 4.183179000976517 | 2.0915895004882583 | 2.0915895004882583 |
| 5 | 5.214101190931181 | 2.6070505954655907 | 2.6070505954655907 |

Differentiating w.r.t. $\theta_i$:

$$\nabla_{\theta_i} \log p_\theta(\tau) = \sum_{t=0}^{\infty} \nabla_{\theta_i} \log \pi_\theta(a_t \mid s_t), \qquad (5)$$

because $\rho$ and $P$ do not depend on $\theta_i$ 5. Use the factorized joint policy:

$$\pi_\theta(a_t \mid s_t) = \prod_{j=1}^{N} \pi_{\theta_j}(a_{j,t} \mid s_t) \; \textit{Therefore} \qquad (6)$$

Therefore

$$\log \pi_\theta(a_t \mid s_t) = \sum_{j=1}^{N} \log \pi_{\theta_j}(a_{j,t} \mid s_t), \qquad (7)$$

and

$$\nabla_{\theta_i} \log \pi_\theta(a_t \mid s_t) = \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} \mid s_t) \qquad (8)$$

6. Plug back into the gradient:

decision-making entity. However, this assumption does not always hold. A multi-agent system is not necessarily better than a monolithic intelligent agent. If the environment permits, it is possible to design one single intelligent agent capable of coordinating everything. Multi-agent systems are particularly effective for tasks whose natural decom- position leads to an explicit distribution of work, or for multi-party coordination problems where the latency of elastic decisions is significant for performance, yet the cost of free use of the services of all the other parties is high. In real-time multi-counterparty derivatives trading among autonomous agents, minimal latency is a requirement; hence, it is preferred to have systems that require minimal trust and that do not impose too many conditions on the participation of the different parties; that is, to introduce the idea of non-governed coordination. However, for many of the participating agents, the latency of decisions im- pacting them is not elastic, given some specific task of either satisfying some regulation, or managing some risk, or just executing a trade itself; in these cases, some natural compound should be formed. Besides simple latency op- timization, other aspects involving traffic intensity should be concerned, like price discovery, latency sensitive-driven schedules, and the security of t

*B. Key Constraints and Considerations in Autonomous Agentic Systems*

The successful development of an autonomous agentic framework that realises the capabilities required for real- time derivatives trading across multiple counterparties

$i$ $\theta$

$\nabla_\theta J(\theta) = E_{\tau\sim\pi}\infty t{=}0$

B

$\nabla_{\theta i} \log \pi_{\theta i} (a_{i,t} \mid s_t)$

D

$R(\tau)$

(9)

must address several important constraints and considerations. The most critical of these is latency;

7. To reduce variance we typically use the return from time

t:

$$G_t = \sum^{\infty}_{k=t} \gamma^{k-t} r_k \qquad (10)$$

Then we can write the standard multi-agent REINFORCE form:

$$\nabla_\theta J(\theta) = E_{\tau\sim\pi} \left[ \sum^{\infty} \nabla_\theta \log \pi_\theta (a_{i,t} \mid s_t) G_t \right]. \qquad (11)$$

meeting the required response times for automated control

requires careful choice of communication mechanisms and protocols for coordinating the agents. In many multi-agent control problems, agents do not initially trust one another. Trust and reputation management is therefore important to ensure that control decisions are based on reliable decisions and information. Derivatives trading often involves different operational systems and procedures across venues and counterparties, and a

$i$ $\theta$ $i$ $i$

$t{=}0$successful agentic framework must, therefore, support

**Simulated Exposure, Margin and Net Risk**

*A. Challenges and Requirements in Autonomous Agentic Systems*

Functionality and reliability considerations are cru- cial to support a self-sufficient orchestration of real-time derivatives trading among multiple counterparties. Besides possessing the conventional qualities of agent-based sys- tems, autonomous agentic systems must also meet spe- cial demands arising from the task of coordinating the operational policies of the agents involved. Multi-agent systems are often deemed suitable for tasks requiring complex decisions based on the global situation, for which distributed intelligence is better than a single monolithic the heterogeneous nature of the settings in which the coordination takes place. The risk of algorithmic or agentic trading disproportionality during stressed markets is difficult to quantify and hence the avoidance of dynamic control oscillations for simple problems does not guarantee the avoidance of similar behaviour for more complex scenarios. The absence of a single trusted venue for execution of orders adds additional challenges in the coordination of the trading strategy. The successful operation of an agentic framework that maps trading strategies to real-time execution orders and timings would indicate feasibility of a multi-party, cross-venue algorithmic trading strategy where the order flows for a product across its complete duration are monitored

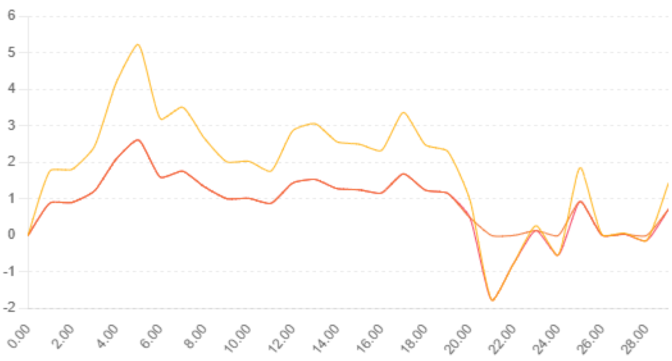


Fig. 2. Stylized Counterparty Exposure and Margin Dynamics

and controlled. The ability to maintain balanced hedges, ensure liquidity and manage counterparty risk across the counterparties on a product and across products would position the trading agents to autonomously manage collateral for the duration of the trades and thereby remain compliant with margin regulations.

**Equation 2 – Real-Time Risk-Aware Reward Shaping**

### 2.1. Basic shaping

Let:

Rt: pure trading P&L at time $\Delta t$.

Ct: collateral or margin cost at time $\Delta t$.

Rt: risk measure at time $\Delta t$ (for instance, a one-step Value- at-Risk).

Define a risk-aware shaped reward:

$$Rt' = Rt - \lambda_C Ct - \lambda_R Rt, \tag{12}$$

where $\lambda_C, \lambda_R \geq 0$ are trade-off coefficients chosen by risk/compliance agents.

The new objective is:

C

formal objectives in a multi-party system are synthe- sized and extended to capture self-coordination and self-adaptation of multiple trading decision-making parties. Second, the description of how a trading context can be modeled as an online learning scenario—with agents receiving real-time streams of information relevant to the interactions and the state of the environment, knowl- edge, or preferences required for the agents to perform successfully being unknown but changing continuously with time, and adaptation cycles occurring much faster than the underlying dynamics—allows for the design of practical self-adaptation mechanisms by defining what should be sensed and monitored, proposed alarms, and appropriate control loops. Third, multi-party order exe- cution strategies describe how a trading order must be factored into sub-orders assigned to specialized executing parties, given that the required mode of execution cannot be decided by any single party in real time due to latency constraints, and that current market conditions, latency with respect to the various possible executing venues, co- movements of correlated instruments, and other aspects of the surrounding environment must be accounted for. Fourth, appropriate risk coverage and margin management techniques identify orchestration actions for ensuring that all market risks (including counterparty risks) are duly covered at all times, by choosing the size of covered positions and margin distributions across counterparties in real time. First, agent-based coordination studies all aspects of a collective intelligence consisting of agents with sufficient autonomy and scope for decision making, with a focus on the patterns of interaction and decision making among the agents. A group of agents is said to act collectively in order to accomplish an objective, usually formal and external to the participating agents, provided they do so through a pattern of interaction (distinct from mere cooperation or simply acting together) that satisfies certain conditions (e.g., is non-redundant, includes agents with different functions, and cannot be reduced to an inter- action among their leaders). These ideas address cooper-

$t$

### 2.2. Effect on the gradient

Because $R'$ is a linear function of ($R$ , $C$ , $R$ ), the policy ation and collaboration—concepts that characterize more

general gatherings of agents taking joint action, whether

$t$ $t$ $t$ $t$ the pattern of interaction is given or not—and extend the

gradient is the same as in Eq. 1 with $R_t$ replaced:ideas of a delegation-supported partnership (used to study

$\nabla_\theta J'(\theta) = \mathrm{E}$ $i$

C$\sum$

$t$

$\nabla_\theta \log \pi_\theta$ (14) $i$ $i$

the joint actions of trusted agents independently pursuing similar objectives) to deal with less-trusted agents.

$(a_{i,t}|s_t)$ $\infty$ A

B

$k{=}t\gamma^{k-t}(R_k - \lambda_C C_k - \lambda_R R_k)$ (15)

*Agent-Based Coordination Principles* Coordination through agents makes many everyday processes easier. Such processes have one or more common

## I. THEORETICAL FOUNDATIONS OF AUTONOMOUS AGENTIC FRAMEWORKS

The foundational theory of autonomous agentic frame- works for real-time multi-party derivatives coordination encompasses four interacting components. First, estab- lished principles of agents acting collectively to accomplish objectives, but also require management of knowledge and information, activity scheduling, negotiation of resources and other multi-party concerns. For example, a household preparing a meal requires one person to coordinate and orchestrate preparations such as cooking, eating and cleaning, while simultaneously organising preparation of any loved ones who arrive late. Although

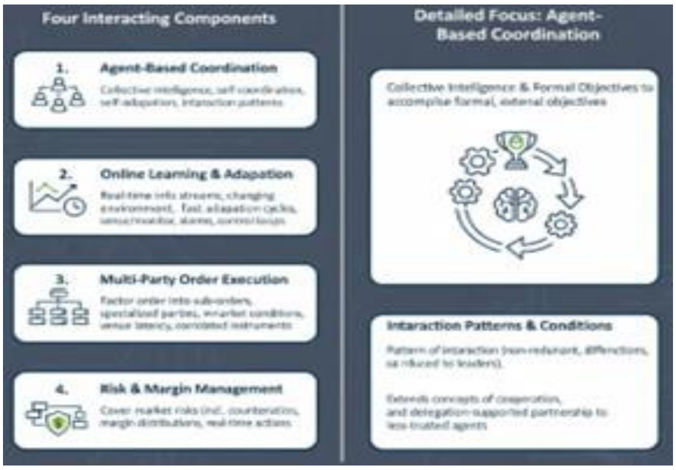


Fig. 3. Autonomous Agents: Foundational Theory for Multi-Party Derivatives Coordination

participating agents may be offered control, any confusion or disagreements ultimately fall responsibility on the coordinating agent. The act of control moves to dependent agents with simpler, ad-hoc tasks. Coordination theory proposes that all systems requiring similar behaviour share a set of common elements that explain both agent roles and their execution. A common set of agent types and behaviour have been used to adapt agent-based coordination into trading, transportation, robotics and other sectors. Despite the current simplicity of most trading systems trading objectives are actually more complex with multiple counterparties interacting with one set of underlying: a derivatives transaction. Many transactions require bespoke, opaque systems based on custom scripting languages, making their creation outside the capabilities of a sole trader. In fact, custom builds are already being generated based on Compiled Code Prototyping, which creates self-contained code by merging multiple calls into a virtual library. An adaptive trading engine would therefore fill a real need for both execution of current trades and as a base upon which new derivatives trading systems could be built.

**Equation 3 – Counterparty Exposure Dynamics**

**3.1. Discrete-time dynamics**

Let:

Et: exposure at time $\Delta t$ (positive means we are owed money).

ΔVt: change in mark-to-market value of the portfolio between $\Delta t$ and $\Delta t + 1$.

Mt: collateral / margin posted to us at time $\Delta t$.

Rt: collateral / margin we posted to the counterparty at time $\Delta t$.

Define net collateral received:

$$Kt = Mt - Rt \qquad (16)$$

A simple linear exposure update is:

$$Et + 1 = (1 - \rho)Et + \Delta V t - Kt \qquad (17)$$

**3.2. Iterated solution**

Unroll the recursion: 1. $E_1 = (1 - \rho)E_0 + \Delta V_0 - K_0$

2. E2=(1-ρ)E1+ΔV1-K1

$$= (1 - \rho)2E0 + (1 - \rho)(\Delta V0 - K0) + (\Delta V1 - K1) \quad (18)$$

By induction, for general t:

$$Et = (1 - \rho)^t E0 + \sum_{k=0}^{t-1} (1 - \rho)^{t-1-k}(\Delta V_k - K_k). \qquad (19)$$

*Self-Adaptation and Online Learning in Trading Envi- ronments*

A trading environment can be viewed as a set of agents that execute key economic functions: informing, trading, and specifying coverage needs associated with a range of risk factors. Learning objectives of these agents support these functions. Usually, the information agents, trading agents, and coverage-specifying agents execute their re- spective functions independently. If an agent observes a significant change in its environment, such as a new trader with different preferences, it will inform other agents of this observation. Based on this information, other agents can react accordingly. The system can be modeled as an online learning framework, where agents learn from the actions of other agents in the past. This learning- adaptation process occurs in cycles. Each cycle consists of observing, sensing, and controlling. Trading agents learn the runtime model and control parameters over these cycles. The backtracking applied by the information agents ensures that the information sharing is resilient to noisy announcements. The communication topology is also significant, especially in the design of nice enough conditions; a nice enough condition for Krasovskii–La Salle theorem ensures that such topologies remove oscillatory dynamics. Such a condition facilitates a shifting-coalition stabilization executive club effect, where agents that are adversely affected by the current action of the coalition can opt out by imitating nearby inactive agents. In such a scenario, if a monitoring agent detects that the population is oscillating between two control actions, it will announce this fact, and one of the actions will be flagged as unsafe. For the existing action that is flagged as unsafe, an agent will also be monitored for possible adversarial behavior.

## II. SYSTEM ARCHITECTURE AND FRAMEWORK OVERVIEW

Core Components and Responsibilities The autonomous agentic framework comprises five core agent types with distinct yet complementary roles: decision, execution, risk, compliance, and reconciliation agents. Decision agents address requests for action—such as trade, lending, or borrowing—by deciding whether to accept, delegate, or reject. Accepting decisions trigger the implementation phase, staged by a designated execution agent. Risk and compliance agents are triggered by detected actions to apply appropriate constraints and controls. Movement of the underlying risk triggers those agents responsible for post-trade reconciliation. The process is illustrated in Figures 2 (operation flow) and 3 (functional roles). A decision agent accepting a request identifies the most appropriate execution agent among those available. The presence of multiple execution agents is essential for effective latency management. The agent coordinating intraparty activity determines the venue(s) where counterparty orders are to be filled; multiple agents can also participate when multiple offer orders on multiple venues are present. The agent

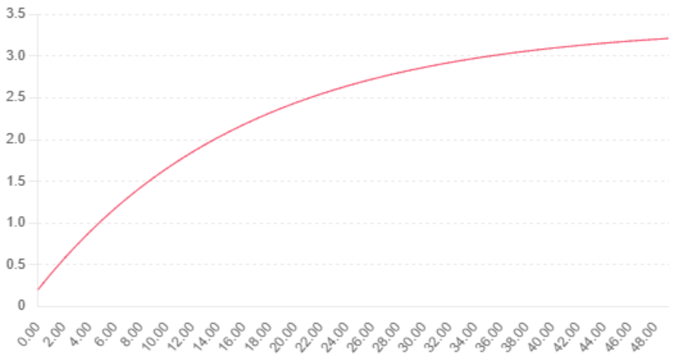


Fig. 4. Illustrative Regret Curve for Online Strategy Adaptation

Define the adaptive objective

considering interparty coordination evaluates whether a lending transaction should occur and informs the

$$\min_{x \geq 0} \sum_{j=1}^{m} c_j x_j + \lambda \sum_{j=1}^{m} \sigma_j(x_j) \text{ s.t. } \sum_{j=1}^{m} x_j \leq B, \qquad (21)$$

respective parties. Communication Protocols and Trust Management Agents communicate through a combination of publish-subscribe messaging via a broker (for responses to requests) and direct messaging (for orders). The messaging protocol is based on one established for financial product disclosures (Futures Industry Association, 2016).with risk-aversion parameter $\lambda \geq 0$.

### 4.2. First-order condition (for unconstrained optimum)

Ignoring the budget constraint temporarily and differenti- ating w.r.t. $x_j$:

Messages cover the entire trading spectrum, are in

$\partial(x \quad + \varepsilon)^2$

JSON format, and support multi-party exchange. Agent identity is based on an account associated with the

$\partial x_j(c \ x \ + \lambda x \ + \varepsilon\beta \ ) = c \ - \lambda \frac{j}{\ }.$ (22)

$j \ j \quad j \quad j \quad j \qquad \beta j$

relevant agent, allowing both formal authentication and domain-based authentication of trades per legal

Set derivative to zero:

$(x_j + \varepsilon)^2$ 2

requirements. Messages may be signed—ensuring sender authenticity, message integrity, and non-repudiation—and$c_j - \lambda \quad = 0 \Longrightarrow (x_j + \varepsilon) = \lambda c_j \beta_j$ (23)

$j$ $\beta$

encrypted, maintaining sender-provided confidentiality (e.g., for proprietary trading strategies). To establish trust, the communication methods in the architecture draw from hierarchical trust models. Dynamic trust is based on context-dependent requirements, specifying how much or how little trust is needed for a particular interaction. Conditions for trust can be negotiated as part of the interaction process. As some agents are intraparty and others interparty, confidentiality and non-repudiation are only applicable for intraparty communication, while message provenance is a priority. Latency management and communication failures—both expected and unexpected—determine whether sensing and machinery redundancies are present.

**Equation 4 – Adaptive Collateral Optimization**

**4.1. Adding risk penalties**

Let $\sigma_j(x_j)$ be a convex risk measure worsened by low collateral (e.g. variance of exposure, tail loss). A simple surrogate is

$\sigma_j(x_j) = x_j + \varepsilon\beta_j$ (20)

where $\beta_j > 0$ and $\varepsilon > 0$ avoids division by zero. More collateral $x_j$ reduces riskTaking the positive root (since $x_j + \varepsilon > 0$):

$x^* = \lambda c_j \beta_j - \varepsilon.$ (24)

$j$

*A. Core Agents and Roles*

Five core agent types realize the self-adaptive multi- agent coordination mechanism permitting real-time pooling of expertise, capabilities, and information from an enterprise's entire resource base. The mechanism can close the coordination gap for trading securities in real-time via multiple market participants or for other applications requiring coordination with multiple counterparties. Such coordination creates new enterprise-scale trading opportu- nities and expands the utility of collateral by facilitating real-time pool management. Joint goal-oriented execution of derivatives trades significantly reduces the costs asso- ciated with coordination among the trade counterparties. Decision agents express preferences defining their favorite decision outcomes with respect to given decision alterna- tives. Execution agents execute a decision outcome accord- ing to a decision agent's specification and execute their own decisions. Risk agents manage risk indirectly through the dynamics of the risk indicators they control. Compli- ance agents make decisions according to their respective goals and preferences and send the resulting commands to the execution agents. Reconciliation agents interact with all major risk management systems providing them with the data necessary for near-real-time risk assessment and detection of reconciliation triggers.

*B. Communication Protocols and Trust Management*

Trust management is essential to the adoption of au- tonomous agents on a large scale. The following aspects must be considered when designing communication be- tween agents managed by different stakeholders: * **Mes- saging standards**: Standard protocols must be used for exchanging messages to minimize integration cost and ensure late discoverability of potential mismatches and errors. The protocols must also permit late binding of message structure and parameter types to allow full in- tegration of unknown agents without requiring schema changes. * **Authentication and confidentiality**: Mes- sages are exchanged across distinct execution contexts. Their structure may change throughout an execution and messages may go through other parties. Authentication mechanisms must be used to guarantee the origin, recip- ients, and integrity of the messages. Proper handling of secrets is necessary to enforce confidentiality and avoid the leakage of business-sensitive information. * **Data provenance**: Information generated from past or future events may be referenced in the execution phase. Spec- ifying its provenance allows agents and the system as a whole to validate information quality before using it to make decisions. * **Latencies**: The latencies implied by the different protocols and operations must be considered. Many protocols have asynchronous semantics, such as request/reply or publish/subscribe. If latency is a critical parameter, the resultant indirect flow of execution must be analyzed and considered in the decision logic. * **Fault tolerance**: A crucial aspect of a real-time execution context is the capability of being resilient to failures. The fault-tolerant logic must be applied to both direct calculation and the responses to unexpected events.

## III. SELF-ADAPTIVE COORDINATION MECHANISMS

Two mechanisms support the self-adapting agentic co- ordination of activities within the proposed framework. The first relies on elaborated goals that drive the decision agents' planning and reactions to context changes. The second entails feedback loops that monitor and control agents' activities and interactions during execution. For- mulating a new goal involves expressing its objectives, preferences, and constraints. Goals are then embedded in hierarchical decision trees, which execution agents oper- ate. Information provided through feedback loops alters the formulation process, assures the execution's stability, and mitigates the risk of oscillations and adverse inter- agent dynamics. Detailed descriptions elucidate these two dimensions of coordination.


Fig. 5. Self-Adapting Agents: Goal-Driven Coordination & Feedback Loops

*1) Goal-Driven Adaptation:* An agent aids a decision agent by expressing its preferences and constraints over the action space or by suggesting a suitable action. Execution agents process these formulations and run decision trees as part of executing their goals. Periodic updates can capture the agents' changing situations as new observations shape their preferences, constraints, and solution feasibility. Sup- porting the decision agents' planning entails allowing other agents to contribute context-specific information. They can express preferences over the venue for executing an order, suggest the latency to reduce for an incoming order, or make recommendations that affect other decision agents' objectives, such as those expressed by a central bank.

*A. Goal-Driven Adaptation*

Goal-driven adaptation, the first self-adaptation mech- anism underpinned by autonomous agentic frameworks, asserts that agents—through higher-level behavioral prin- ciples and internal knowledge bases—must possess the autonomy to modify their activities in light of changing objectives and circumstances. Fulfillment relies on the formulation of goals (or goals for the agents to act on, as conceptually distinct from targets for agent coordina- tion) through the interrogation of decision agents based on their immediate objectives, preferences, constraints, and sensitivities to other agents' operations. Goals can range in scope from localized, immediate-level goals—only acting on, or being acted on by, a small number of agents within immediate proximity or latency requirements—to increasingly more ambitious objectives that traverse in- creasing numbers of agents, actor types and horizons. Where temporal hierarchies exist, high-level goals set overall directions but may be fully subsumed by low-level objectives. Two scenarios for the goal-driven modification of agent activities are worth illustrating. In an illustrative trading application, the embedding of a crucial configura-

**3. Aggregated communication vector:**

tion parameter in an agent's behavioral principle allows intercept signals to it $\sum$ to probe "what market opportunity would have to change (e.g. what trading price would have $c_{k,t}$ =

$i \in N(k)$

$\alpha_{ki,t}$

$z_{i,t}$. (27)

to drift farther, what inventory constraint would have to become tighter, or what risk would have to grow bigger) to drive a change of direction in action?" Alerts returned by the agent allow other agents to sense the proximity of such turning points. Other nearby actors may—if configu- ration states permit—re-schedule objectives to establish co-ordinated trade within those parameters, whilst still ensuring a full set of delaying or proxy trades. In yet another application, tracers across a duration horizon can feed-back control top-down toward two complex, agent- embedded climate profiles, one for natural supply on a price-sensitive basis, and the other for eventual costly, adverse-response-adjusted storage retreat, sensitive also to excess coolant temperature.

*B. Feedback Loops and Stability Considerations*

Adaptive coordination extends beyond goal preference changes, employing feedback loops to monitor the adaptation process itself. To ensure stability, several aspects require attention. First, decisions should minimize state or objective function value oscillations. Second, control parameters need adequate tuning to guarantee closed-loop stability; while the convergence of discrete interactions suffices for stability, autonomous agentic systems enable both continuous and sampled data control. Sensing, monitoring, and alarm conditions are specified for each feedback control loop. Each agent observes a subset of its neighbors, and control actions are initiated if alarm conditions are met, prompting control parameter adjustment. These agents act as sensors with reduced dynamics, addressing trend or result deviations with low effort. Control loops can be layered or modularized. For instance, goal formation and preference processing form a closed control loop within the third- level hierarchy. In trading activities, agents sense trading patterns and offer warnings if persistent adversarial relations or price manipulations arise. Oscillations or shifts in trading patterns indicate potential problems.

**Equation 5 – Inter-Agent Communication En- coding**

**5.1. Attention-based aggregation**

For agent $k$ receiving messages from neighbors $i \in N(k)$, construct:

**1. Compatibility scores:**

$e_{ki,t} = g_\psi(h_{k,t-1}, z_{i,t})$ (25)

where $h_{k,t-1}$ is agent $k$'s previou hidden state, and $g_\psi$ is a scoring network. **2. Attention weights (softmax):**

Agent $k$ updates its internal state:

$h_{k,t} = RNN_\omega(h_{k,t-1}, c_{k,t}, s_t)$, (28) which then feeds into the policy $\pi_\theta\ (a_{k,t} \mid h_{k,t})$.

## IV. MULTI-PARTY ORCHESTRATION FOR DERIVATIVES TRADING

$k$

Multi-party orchestration mechanisms for enterprise- scale derivatives trading have two main components: order execution across multiple trading venues or platforms, and risk coverage across multiple financial counterparties. Orchestrating order execution involves determining suitable execution venues and strategies and ensuring that their specific constraints are respected. Routing rules may incorporate considerations such as transaction costs, venue liquidity, latency in trade registration, and experience with payment delays. Latency in trade registration is particularly important for capital markets such as futures and options, in which an active and latent market exists for very short durations. Orders should preferably be submitted to the venue where they are expected to receive the fastest execution, although this may not be the venue offering the most favorable price. An early venue execution may also influence subsequent trades in other venues, and active arbitrage may lead to pre-hedging against positions taken in markets with lower latencies. Typical venue-specific constraints also include rules prohibiting the disclosure of a trader's identity before trade execution. Risk coverage and margin management deal with the multiparty aspect of derivatives trading. A multi-counterparty transaction exposes its participants to multiple predefined risks, and these risks need to be covered in real time. In addition to managing the risk that the counterparty of the next execution may default, traders need to consider margin requirements when opening or closing a net position. When executing multiple orders in sequence rather than one large order, liquidity requirements must be anticipated and, if possible, planned in order to avoid resorting to costly alternatives such as establishing an overdraft. Moreover, in the case of clearinghouse collateral, it is also important to match the counterparty providing the cheapest collateral for each trade, since margin requirements can vary significantly from one counterparty to another, or become negative. Finally, regulatory requirements must be met to be able to execute even transactions satisfying all the other requirements.

**Equation 6 – Online Strategy Adaptation via Regret Minimization**

$j \in N(k)$ $\exp(e_{kj,t})$

$\exp(e_{ki,t})$(26)

6.1. Definition of regret

Consider a set of possible strategies $H$ (e.g. parameterized

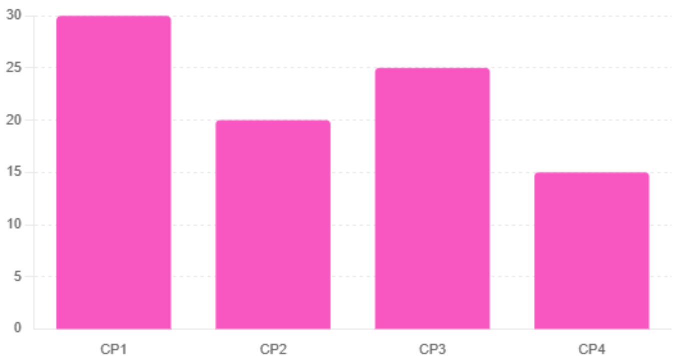


Fig. 6. Example Optimal Collateral Allocation Across Counterpar- ties

policies, or a finite set of hedging rules). At each time $t$: The algorithm picks strategy $h_t \in H$.

Nature reveals a loss function $\ell_t : H \rightarrow$ R (e.g. negative

Sharpe, realized risk, etc.).

Algorithm suffers loss $\ell_t(h_t)$.

The cumulative regret after T steps is defined as:

many derivatives contracts must remain hedged during their lifetime to avoid huge losses should the real world deviate from the expectations over the risk-free interest curve. Can the derivatives traders cover the counter-party risk and the resulting margins at the same time? The margin requirements for non-cleared over-the-counter derivatives can represent a significant share of a dealer's balance sheet; making this management task even more critical for the dealer is the fact that liquidity is also a major driver of funding costs. Collateral is increasingly seen as a source of liquidity that can be optimized, and the management of counterparty margins both improves liquidity and supports netting of risk. Last but not least, bank regulators are also asking their supervised dealers to disclose enough information on their risk management and pricing, so that market participants can evaluate not only the counterparty, but also the counterparty risk that is associated with clearing.

*B. Risk Coverage and Margin Management*

Exploiting the business-critical need to balance risk and return for each derivatives transaction across mul-

$$RT = \sum_{t=1}^{T} \ell_t(h_t) - \min_{h \in H} \sum_{t=1}^{T} \ell_t(h). \quad (29)$$

tiple counterparty relations is one of the largest chal- lenges in an autonomous (agent-based) framework for real-time inter-counterparty derivatives operations. Managing

Desideratum: $R_T = o(T)$ (sub-linear), meaning average regret $R_T/T \rightarrow 0$. **6.2. Gradient-based updates**

If H is parameterized by $\theta$ and each $\ell_t(\theta)$ is differentiable:

Online gradient descent:

$$\theta_{t+1} = \theta_t - \eta_t \nabla_\theta \ell_t(\theta_t) \quad (30)$$

Under standard conditions (convexity, suitable learning rates), we have theoretical bounds like

$$RT \leq O(T), \quad (31)$$

which implies average regret →0→0.

*A. Order Execution Strategies Across Venues*

When trading orders across multiple counterparties and trading venues, decision agents can achieve performance benefits by dynamically selecting the most appropriate venue for an order depending on the evolving market conditions. Such a determination can balance competing considerations, such as successfully executing the order on-execution price and minimizing relative latency. Addi- tionally, orders for the same derivative instrument placed on different venues can be used for market-making or hedging purposes. The execution of these orders must satisfy all retail and brokerage clients' preferences and relevant market rules, in particular in what concerns order display, time priority and venue-specific tick sizes. Orders sent to different venues are connected through price- business logic: the model must ensure that executing an order on one trading venue will not create and unhedged risk exposure; execution logic over the control plane must transfer shares to avoid arbitrage profits. Furthermore,

margins—which include clearing fund deposits and other forms of collateral—across clearing establishments, exe- cuting dealers, and risk-coverage managers to maximise obsolescence with respect to posted margin represents a major challenge. Templates for risk cover-related margin postings that provide the required breadth and depth of covers should be orchestrated. Additional risk-cover devices not present in order specifications should be con- tinuously monitored during order execution and harnessed for monitoring if such further cover assets become available during potentially protracted order placements. Ordering and covering strategy formulation must be paired together atomically with respect to each margin-updating time step and taken in a way that captures not just the actual apex of margin consumption for the present position but the full future trajectory, where sufficient redundancy sys- tems must exist bilaterally to permit such risk-operating manoeuvring for different operators and hedgers and to bolster real-time latency requirements. These diverse re- quirements combine to stimulate a process of monitoring margin consumption over time on all positions and de- mand that it be catered for as best as possible in real- time order placement, rather than merely at the closure of positions.

## V. CONCLUSION

Real-time multi-counterparty orchestration for deriva- tives trading with autonomously driven agents. The enterprise-size impacts of these novel coordination princi- ples make a compelling case for their resolution of key dif- ficulties in trading and margin management across venues

via an accessible goal formulation process. The second proposed an internal feedback mechanism for adapting learning with respect to changes and alarm signals re- ceived from other parts of the framework. Achieving jointly beneficial multi-party interactions with poorly aligned or untrusted parties through a non-agentic approach would be particularly hard. Nevertheless, together, the two self- adaptive mechanisms should be safer forms of adaptation than candidate adaptations proposed in the literature that assumed a stronger notion of self-stability. Future research could include extending the current architecture with the capability of managing counterparty risk price discovery, latency-aware order submission and margin management with collateral optimisation across the entire trading port- folio. Such complementary capabilities should further en- hance the support for real-time derivatives orchestration across multiple counterparties and widen the potential application scope.

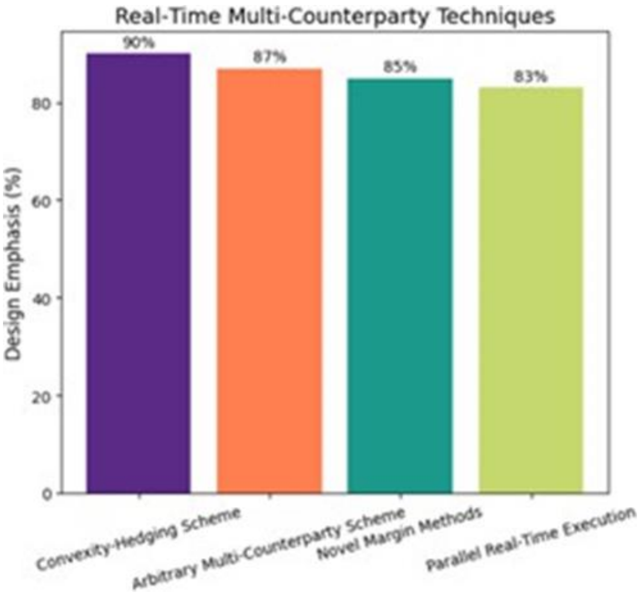


Fig. 7. Real-Time Multi-Counterparty Techniques

with multiple counterparties. These principles certainly warrant further investigation and refinement, with more specialised aspects of the problem domain tackled by dedicated entities, as may be found in the boundaries of associated literatures, such as trust, risk and margin management. Two real-time multi-counterparty function- ing techniques in derivative dimensions, one based on hedging convexity, the other fundamentally arbitrary, and deployable in novel margin management methods, have also been identified. Parallel real-time execution appeals. Yet such trading is non-trivial, presenting challenges of liquidity, price discovery and latency. On the latter, hedg- ing remains an open question, one that might leverage advanced trading contemplations, such as market-making ghost orders placed in external events. Until addressed by available methodologies, trading—perhaps opportunity trading—will need to be further explored, but from the point perspective of optimally episodic or at least reduced- order spectral re-sampling.

### A. Final Reflections and Future Directions

The research presented a framework facilitating self- adaptive and agent-based coordination capabilities. The defined key principles enable novel support for real-time derivatives transactions involving multiple counterparties interacting through heterogeneous protocols with poten- tially different collateralisations. The trading and orches- tration agent layers provide a multi-party view not achiev- able through individual agents residing within the coun- terparties. Two high-level mechanisms were introduced, realising flexibility in self–adaptation without periodic retraining or explicit modelling of harmful dynamics. The first discussed how learning direction, preferences, exe- cution strategy or coordination policies can be updated

REFERENCES

[1] Abreu, D., & Sannikov, Y. (2022). A theory of dynamic con- tracting. *Econometrica, 90*(5), 1931–1974.

[2] Adrian, T., Boyarchenko, N., & Giannone, D. (2019). Vulnerable growth. *American Economic Review, 109*(4), 1263–1289.

[3] adi, A. L. The Role Of AI-Driven Predictive Analytics In Automotive R&D: Enhancing Vehicle Performance And Safety.

[4] Ahn, D.-H., & Gao, J. (2020). A theory of hedge fund risk premia. *Review of Financial Studies, 33*(2), 708–742.

[5] Athey, S., & Imbens, G. (2019). Machine learning methods that economists should know about. *Annual Review of Economics, 11*, 685–725.

[6] Balduzzi, D., & Ortiz, C. (2022). Multi-agent coordination and learning dynamics. *Journal of Machine Learning Research, 23*(1), 1–49.

[7] ahari Pandiri, "Leveraging AI and Machine Learning for Dy- namic Risk Assessment in Auto and Property Insurance Mar- kets," International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJIREE-

ICE), DOI 10.17148/IJIREEICE.2023.111212

[8] Bichler, M., Calvano, E., & Lankford, M. (2021). Market design for algorithmic trading environments. *Operations Research, 69*(6), 1792–1815.

[9] Bichler, M., & Füllbrunn, S. (2022). Price discovery in high- frequency markets. *Management Science, 68*(1), 55–75.

[10] Brunnermeier, M. K., & Oehmke, M. (2013). The maturity rat race. *Journal of Finance, 68*(2), 483–521.

[11] Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithms, and collusion. *American Economic Review, 110*(10), 3267–3297.

[12] Cartea, Á., Jaimungal, S., & Penalva, J. (2015). *Algorithmic and high-frequency trading*. Cambridge University Press.

[13] andan, B. P., & Chitta, S. S. (2023). Machine Learning Driven Metrology and Defect Detection in Extreme Ultraviolet (EUV) Lithography: A Paradigm Shift in Semiconductor Manufactur- ing. Educational Administration: Theory and Practice, 29 (4), 4555–4568.

[14] Chen, Y., Wen, Q., & Li, X. (2021). Reinforcement learning in multi-agent financial systems. *IEEE Transactions on Neural Networks and Learning Systems, 32*(9), 3892–3906.

[15] Cont, R., & Kan, Y. H. (2019). Machine learning for dynamic portfolio selection. *Quantitative Finance, 19*(6), 1037–1053.

[16] Cont, R., & Kokholm, T. (2014). Central clearing of OTC derivatives. *Operations Research, 62*(2), 239–265.

[17] oppolu, H. K. R., Sheelam, G. K., & Komaragiri, V. B. (2023). Autonomous Telecommunication Networks: The Convergence of Agentic AI and AI-Optimized Hardware. International Journal of Science and Research (IJSR), 12(12), 2253-2270.

[18] Cunningham, T., & Lacity, M. (2022). Autonomous intelligent agents in financial markets. *Journal of Information Technology, 37*(4), 409–431.

[19] Deissenberg, C., van der Hoog, S., & Hommes, C. (2008). A multi-agent model of the financial system. *Journal of Economic Dynamics & Control, 32*, 361–378.

[20] alisetty, S., & Singireddy, J. (2023). Agentic AI in Retail: A Paradigm Shift in Autonomous Customer Interaction and Supply Chain Automation. American Advanced Journal for Emerging Disciplinaries (AAJED) ISSN: 3067-4190, 1(1).

[21] Duffie, D., & Zhu, H. (2011). Does a central clearing coun- terparty reduce counterparty risk? *Review of Asset Pricing Studies, 1*(1), 74–95.

[22] Easley, D., López de Prado, M., & O'Hara, M. (2012). Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies, 25*(5), 1457–1497.

[23] Foucault, T., Pagano, M., & Röell, A. (2013). *Market liquidity: Theory, evidence, & policy*. Oxford University Press.

[24] Ganesh, S., Giesecke, K., & Kim, B. (2023). Risk-sensitive optimal trading and execution. *Mathematics of Operations Research, 48*(2), 597–618.

[25] akkarasu, P. (2023). Generative AI in Financial Intelligence: Unraveling its Potential in Risk Assessment and Compliance. International Journal of Finance (IJFIN)-ABDC Journal Qual- ity List, 36(6), 241-273.

[26] Gomber, P., Haferkorn, M., Lutat, M., & Zimmermann, K. (2021). Algorithmic trading compliance and supervision. *Jour- nal of Trading, 16*(1), 7–23.

[27] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learn- ing*. MIT Press.

[28] Hamid, K., & Heiden, M. (2022). Agent-based modelling of systemic risk. *Journal of Economic Behavior & Organization, 203*, 119–140.

[29] ummari, D. N. (2023). Energy Consumption Optimization in Smart Factories Using AI-Based Analytics: Evidence from Auto- motive Plants. Journal for Reattach Therapy and Development Diversities. https://doi. org/10.53555/jrtdd. v6i10s (2), 3572.

[30] Hommes, C. (2021). Behavioral and agent-based macrofinance.

*Journal of Economic Literature, 59*(1), 149–219.

[31] Hendershott, T., & Riordan, R. (2013). Algorithmic trading and market liquidity. *Journal of Financial and Quantitative Analysis, 48*(4), 1001–1024.

[32] heelam, G. K. (2023). Adaptive AI Workflows for Edge-to-Cloud Processing in Decentralized Mobile Infrastructure. Journal for Reattach Therapy and Development Diversities. https://doi. org/10.53555/jrtdd. v6i10s (2). 3570ugh Predictive Intelligence.

[33] Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance, 34*, 2767–2787.

[34] Kroll, J. A., & Narayanan, A. (2022). Designing accountable algorithms for financial systems. *Communications of the ACM, 65*(2), 44–53.

[35] LeBaron, B. (2023). Computational finance and agent-based models. *Journal of Economic Dynamics & Control, 147*, 104549.

[36] otamary, S. (2023). Integrating Intelligent BSS Solutions with Edge AI for Real-Time Retail Insights and Analytics. European Advanced Journal for Science & Engineering (EAJSE)-p-ISSN 3050-9696 en e-ISSN 3050-970X, 1(1).

[37] Li, Y., & Wang, S. (2020). Multi-agent reinforcement learning for dynamic trading. *Expert Systems with Applications, 152*, 113365.

[38] Macrina, A., & Mahomed, O. (2020). Funding valuation adjust- ment: Theory and computation. *Quantitative Finance, 20*(7), 1183–1201.

[39] Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets, 16*(4), 712–740.

[40] eda, R. (2023). Data Engineering Architectures for Scalable AI in Paint Manufacturing Operations. European Data Science Journal (EDSJ) p-ISSN 3050-9572 en e-ISSN 3050-9580, 1(1).

[41] Min, H., & Mallipeddi, R. (2021). Multi-agent collaboration in financial decision automation. *Decision Support Systems, 143*, 113494.Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human- level control through deep reinforcement learning. *Nature, 518*(7540), 529–533.

[42] omu, B. (2023). Towards Self-Healing Bank IT Systems: The Emergence of Agentic AI in Infrastructure Monitoring and Management. American Advanced Journal for Emerging Dis- ciplinaries (AAJED) ISSN: 3067-4190, 1(1).

[43] O'Hara, M. (2015). *High-frequency market microstructure*. Oxford University Press.

[44] Schär, F., & Berentsen, A. (2020). *Blockchain and the future of financial infrastructure*. MIT Press.

[45] nala, R. Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspec- tive.

[46] Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities. *Future Generation Computer Systems, 108*, 244–258.