finiz
Singidunum University International Scientific Conference
Međunarodna naučna konferencija Univerziteta Singidunum

PAPERS FROM THE THEMATIC AREAS OF THE CONFERENCE
Corporate Finance
Original paper

MULTI-AGENT AI FOR ADAPTIVE TREASURY AND CAPITAL OPTIMIZATION

Marko Mihajlović[1]*,
Angelina Njeguš[2]

[1]PhD candidate,
Singidunum University,
Belgrade, Republic of Serbia

[2]Singidunum University,
Belgrade, Republic of Serbia

Abstract:

Enterprises face growing challenges in managing liquidity and capital flows amid market volatility, regulatory complexity, and fragmented digital infrastructures. Recent industry research, including The Harmony Gap study by FIS and Oxford Economics, reveals that organizations incur nearly $100 million in annual losses due to inefficiencies such as cyber threats, fraud, and operational friction. This paper examines these challenges through the lens of business agility and introduces a multi-agent artificial intelligence (AI) framework for adaptive treasury and capital optimization. Using a modular, agent-based design, the system combines forecasting, liquidity management, FX (foreign-exchange) risk monitoring, and instrument recommendations to enhance visibility, resilience, and efficiency. By integrating data-driven forecasting with agent-based decision-making and rules-based safeguards, the framework aims to reduce idle balances, maximize yields, and improve compliance. We explore how multi-agent AI can transform treasury operations, enabling organizations to achieve agility, mitigate risks, and remain competitive in dynamic global markets.

Keywords:

multi-agent AI, treasury management, capital optimization,liquidity management, business agility.

1. INTRODUCTION

In today's global economy, enterprises must navigate unprecedented uncertainty caused by volatile financial markets, complex regulations, rapid digital change, and increasing competition. Business agility has become a key factor for sustainable growth, allowing organizations to respond quickly to shifting market conditions while staying resilient and efficient (Kocot, 2023). However, recent studies show ongoing inefficiencies that reduce enterprise value. The Harmony Gap study, conducted by FIS and Oxford Economics (2025) indicates that large organizations lose nearly $100 million each year due to disharmony across the financial lifecycle, with the most friction happening when money is in motion. Significant challenges include idle liquidity, fragmented treasury systems, higher exposure to FX risk, and operational inefficiencies from manual processes.

Despite a growing academic interest in multi-agent systems, existing approaches rarely address enterprise treasury holistically. Current solutions typically focus on isolated aspects, such as forecasting, FX trading, or liquidity, but lack an integrated, adaptive framework for treasury operations. This creates a clear gap: enterprises need systems that unify liquidity forecasting, FX risk monitoring, and instrument recommendation under a single, intelligent architecture.

Correspondence:
Marko Mihajlović

e-mail:
marko.mihajlovic.23@singimail.rs

DOI: 10.15308/finiz-2025-78-87

This paper introduces a multi-agent AI framework for adaptive treasury and capital planning. The solution employs modular, agent-based intelligence to automate liquidity management, address cross-currency risk, and stay compliant with evolving regulations. Compared to traditional treasury management systems, which rely heavily on manual oversight, static forecasts, and disconnected infrastructure, this framework provides adaptive, rule-based decision-making that enhances visibility, resilience, and efficiency.

The contributions of this work are:

1. Improving visibility and transparency across fragmented treasury systems.
2. Reducing idle balances and increasing yield through adaptive liquidity deployment.
3. Incorporating rule-based compliance for evolving regulations in payments, FX, and digital assets.
4. Showing business agility by enabling treasury functions to respond dynamically to market changes.

The introduction of multi-agent AI systems into treasury management carries certain risks. Implementing such systems involves addressing issues related to data quality, cybersecurity, integration costs, and regulatory compliance (Leitner *et al.*, 2024). Nonetheless, the potential benefits, such as optimized yields, reduced operational inefficiencies, real-time auditability, and enhanced business agility, make this an increasingly important innovation (Shirzad & Rahmani, 2024). By aligning with emerging regulations on digital assets, payments, and FX disclosures, the framework provides a forward-looking model for enterprises aiming to succeed in a volatile financial environment.

2. LITERATURE REVIEW

The connection between AI and organizational agility is increasingly being examined. Atienza-Barba *et al.* (2024) show that generative AI greatly enhances agility by supporting responsive decision-making in uncertain situations. Rastogi & Pandita (2025) demonstrate that AI-driven transformation, along with workforce agility, builds dynamic capabilities crucial for entrepreneurial resilience. These studies highlight that AI isn't just about automation; it's a strategic tool for adaptability. (Fosso Wamba *et al.* (2022)) Additionally, confirm that AI integration enhances firm performance by increasing organizational agility.

AI and machine learning have shown potential in treasury applications, especially for improving cash forecasting. Donepudi *et al.* (2020) systematically review ML's role in predicting cash flows and support centralized treasury systems to enable effective AI use. Furthermore, the broader AI revolution in finance, covering fraud detection, risk management, and credit scoring, has been widely documented, though issues related to explainability, systemic risk, and ethics still remain (Maple *et al.*, 2023).

Multi-agent architectures are increasingly popular for dynamic and decentralized decision-making. Ayatollahi & Jafari (2025) introduce Orchestrated Intelligence, a multi-agent system designed for high-frequency quantitative risk management in financial trading. Similarly, Hernes *et al.* (2024) present A-Trader, a multi-agent platform that supports forex decision-making through technical, fundamental, and behavior-based analyses. Workflow optimization using multi-agent systems has also been explored; for example, Dodnikov (2024) shows how MAS can streamline financial operations, though scalability and agent interoperability still pose significant challenges.

These components, including AI-enabled agility, AI in treasury, and MAS in finance, form a strong theoretical foundation for the proposed multi-agent AI framework. They highlight the importance of adaptive, autonomous systems in boosting resilience, liquidity, and strategic alignment during changing conditions (X. Wang *et al.*, 2022). However, despite progress in AI-driven agility and multi-agent designs for trading and risk management, few solutions specifically address enterprise treasury operations with an integrated approach. This paper fills that gap by integrating forecasting, compliance, and decision automation into a modular, agent-based system, shifting from reactive AI to proactive, agility-enhancing intelligence.

3. METHODOLOGY

This study adopts a dual-method approach, combining secondary data analysis of The Harmony Gap study (FIS & Oxford Economics, 2025) to ground the problem context with the conceptual design of a multi-agent AI framework for adaptive treasury and capital optimization.

3.1. RESEARCH OVERVIEW

This study builds on the Harmony Gap research conducted jointly by FIS and Oxford Economics in 2024–2025. The project consisted of two large-scale surveys of senior executives and business leaders directly involved in financial technology decision-making across the United States (45%), the United Kingdom (45%), and Singapore (10%). Each survey included 501 respondents, resulting in a combined dataset of over 1,000 participants from both financial institutions (60%), including banking, asset management, private equity, and fund administration, and non-financial sectors (40%) such as technology, fintech, and insurance.

The first "pulse survey," carried out in October–November 2024, examined sources of financial disharmony, including fraud, cyberthreats, human error, operational inefficiencies, and regulatory complexity, while also highlighting potential growth opportunities. The second survey, in November–December 2024, focused on how organizations are implementing strategies to mitigate these inefficiencies. Data were collected through a combination of computer-assisted telephone interviewing (CATI) and online methodologies, ensuring both depth and comparability across industries.

FIS, a Fortune 500 financial technology leader, provided industry expertise and access to organizations across the global money lifecycle. At the same time, Oxford Economics, a leading independent advisory firm, ensured methodological rigor and analytical depth. Together, their collaboration produced one of the most comprehensive assessments of inefficiencies in enterprise financial operations to date, revealing that large organizations lose an average of $98.5 million annually due to disharmony in their money lifecycle.

3.2. SECONDARY DATA ANALYSIS

This paper uses a secondary data analysis approach to explore the financial and operational issues in enterprise treasury management. Secondary analysis enables researchers to gain new insights from existing high-quality datasets, thereby avoiding the costs and limitations associated with collecting primary data (Johnston, 2014).

The survey results show that organizations lose an average of $98.5 million each year because of issues in the money lifecycle. The most considerable costs are associated with cyber threats ($31.7M), fraud ($21.6M), regulatory complexity ($14.9M), and operational inefficiencies ($11M). More than half of respondents pointed out that the money-in-motion phase, covering payments, transfers, and settlement processes, has the most friction. These findings help shape the problem framing of this paper by highlighting where treasury inefficiencies are most severe.

The secondary analysis in this research does not attempt to replicate or rerun statistical models from the FIS/Oxford dataset. Instead, it synthesizes reported findings to:

1. Determine the scope and importance of the issue.
2. Emphasize persistent inefficiencies in liquidity management, FX risk, and compliance despite digital transformation.
3. Identify technological gaps where advanced AI and multi-agent systems could provide value.

This evidence base provides a clear rationale for developing the proposed multi-agent AI framework as a future-oriented solution to the documented inefficiencies.

3.3. FRAMEWORK DESIGN

Building on the documented inefficiencies in The Harmony Gap study and insights from the literature, we propose a multi-agent AI framework for adaptive treasury and capital optimization. The framework is structured around modular, domain-specific agents coordinated by a central recommendation engine. Each agent is responsible for a specialized domain task but operates under rule-based safeguards to ensure compliance, resilience, and business agility.

3.3.1. Agent Roles and Guardrail Rules

The Liquidity Agent ensures efficient use of cash across multiple accounts, subsidiaries, and currencies. It aims to minimize idle balances, avoid costly overdrafts, and optimize working capital. By forecasting incoming and outgoing flows, the agent determines where cash surpluses should be redeployed and where deficits must be covered.

- Sample Rule: Transfer cash on T+1 to balance accounts, ensuring no account reflects a negative balance, since overdrafts incur penalty charges (e.g., 9% interest).
- Operational Guardrails:
  - Cash can only be moved between accounts belonging to the same enterprise group (e.g., from surplus subsidiaries to deficit ones).
  - Transfers are prioritized toward accounts forecasted to experience negative balances within the next 24–48 hours.

Beneficial because it reduces borrowing costs, minimizes external credit line usage, and improves overall liquidity efficiency.

The FX Agent monitors exchange rate exposures across currencies and initiates hedging actions to protect the firm from unfavorable swings. It balances cost-efficiency with risk coverage by determining when exposure thresholds justify hedging.

- Sample Rule: Execute hedging only when projected currency volatility exceeds ±2% vs. baseline, avoiding unnecessary trades and fees.
- Operational Guardrails:
  - FX hedges can only be executed through pre-approved instruments (e.g., forwards, swaps) and counterparties.
  - Exposure calculations are netted at the group level to prevent redundant trades.

Beneficial because it reduces earnings volatility, shields against unexpected currency losses, and avoids over-hedging costs.

The Interest Rate Agent integrates forward-looking yield curve data (e.g., SOFR, EURIBOR) to optimize cash allocation strategies. It helps balance short-term liquidity needs with long-term yield optimization.

- Sample Rule: If forward yield curves project falling rates, prioritize allocating balances to short-term instruments (e.g., T-bills); otherwise, ladder maturities to capture higher yields.
- Operational Guardrails:
  - Investments can only be made in instruments within approved credit ratings and maturity bands.
  - Allocation is capped at defined percentages of total treasury reserves to maintain liquidity buffers.

Beneficial because it prevents misallocation of funds during rate transitions and ensures compliance with investment policy limits.

The Market Sentiment Agent incorporates external indicators, such as volatility indexes, credit spreads, and macroeconomic news, to provide a forward-looking adjustment to treasury decisions. It prevents strategies from being too static in volatile environments.

- Sample Rule: Adjust exposure levels only when sentiment indicators (e.g., VIX > 25, credit spreads widen > 50 bps) cross critical thresholds, avoiding reaction to short-term noise.
- Operational Guardrails:
  - Sentiment signals are weighted alongside fundamental data (cash flows, forecasts) to avoid over-reliance.
  - Adjustments cannot exceed defined tolerance bands (e.g., +/- 10% shift in allocations).

Beneficial because it helps enterprises remain proactive, enhancing resilience without succumbing to market overreactions.

The Recommendation Engine consolidates all agent outputs into a unified treasury strategy, ensuring consistency, compliance, and alignment with enterprise policies. It applies guardrails to enforce regulations, audit requirements, and board-level investment guidelines.

- Sample Rule: Consolidate all agent outputs into a unified strategy, but override recommendations if they violate compliance constraints, disclosure requirements, or internal policy limits.

- Operational Guardrails:
  - o All recommendations pass through a compliance filter (e.g., MiFID II, Dodd-Frank, local FX reporting).
  - o Execution is restricted to counterparties, instruments, and payment rails pre-approved by the enterprise's treasury, risk, or compliance teams.

Beneficial because it ensures that automated intelligence remains explainable, auditable, and aligned with enterprise governance, reducing regulatory and reputational risk.

Figure 1. Proposed Multi-Agent AI Framework for Adaptive Treasury

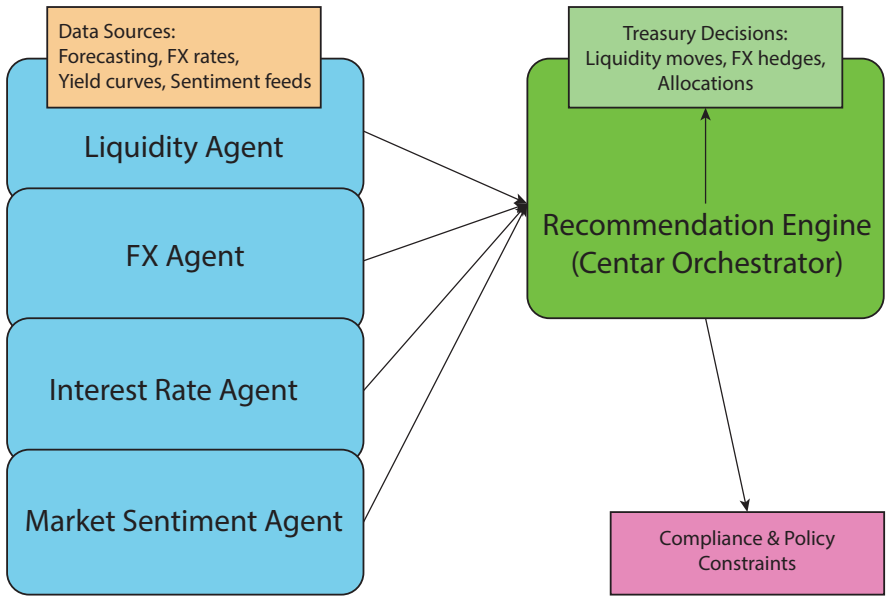

Source: Author-generated illustration

Figure 1 illustrates the modular agents (liquidity, FX, interest rate, and market sentiment) that feed into a central recommendation engine, which applies compliance guardrails before producing unified treasury decisions.

Table 1. Harmony Gap Constructs and Framework Mapping

Harmony Gap Construct (Evidence)	Treasury/Capital Pain Point	Framework Response	Responsible Agent(s) / Controls
Money-in-motion is the top friction (51%)	Settlement timing, routing, and intra-day liquidity	Real-time liquidity mobility, automated transfers, guardrails to avoid overdrafts and fees	Liquidity Agent with money-movement rules
Cyber, fraud, and regulatory are the most significant cost drivers	Losses, audit burden, compliance complexity	Rule-based guardrails, automated logging, policy-aware decision filters	Recommendation Engine + compliance layer
Operational inefficiency & human error persist	Manual cash moves, forecast errors	Automated forecasting, execution safeguards, and exception handling	Liquidity Agent + forecasting models; Recommendation Engine
AI adoption is high, but hindered by cost/skills/ integration	Change risk, limited scalability	Modular, API-first, separation of models from business logic	Microservices, CrewAI orchestration, plug-in agents
Embedded-finance yield signal (~8.5% sales growth)	Value capture from real-time orchestration	Cross-rail capital allocation across MMFs, T-bills, repos, stablecoins	Recommendation Engine consolidating all agents

Source: The Harmony Gap Report, FIS & Oxford Economics (2025)

Table 1 provides a structured mapping between empirical findings from The Harmony Gap study (FIS & Oxford Economics, 2025) and the design requirements of the proposed framework. The mapping highlights how specific pain points, such as money-in-motion friction or regulatory cost drivers, translate into technical features, agent responsibilities, and guardrail controls.

3.3.2. Implementation Options for Agents and Orchestrator

While the framework is conceptual, multiple implementation pathways are possible depending on enterprise infrastructure, cost, and regulatory requirements. In addition to deployment architectures, agent design can draw on established AI techniques tailored to specific treasury tasks:

- Liquidity Agent: Time-series forecasting models such as ARIMA, LSTM networks, or gradient boosting (e.g., XGBoost, LightGBM) can predict cash inflows and outflows.
- FX Agent: Reinforcement learning (RL) methods, including Deep Q-Networks (DQN) and policy gradient algorithms, are well-suited for dynamic hedging under uncertainty.
- Interest Rate Agent: Yield-curve forecasting can employ recurrent neural networks or transformer-based models to optimize allocation strategies.
- Market Sentiment Agent: Pre-trained natural language processing (NLP) models (e.g., BERT-based sentiment classifiers) can analyze financial news, volatility indices, and credit spreads.

These models would be trained on historical financial data and continuously updated with real-time inputs, enabling agents to balance predictive accuracy with adaptability. Prior multi-agent systems in finance, such as A-Trader for forex trading (Hernes *et al.*, 2024) and orchestrated intelligence for quantitative risk management (Ayatollahi & Jafari, 2025), demonstrate the feasibility of such approaches.

For orchestration and deployment, several options exist:

- Microservices with Container Orchestration: Agents deployed as Docker containers managed via Kubernetes or OpenShift, with the orchestrator as an API gateway.
- Serverless Functions: Event-driven execution with cloud-native services (Azure Durable Functions, AWS Step Functions, Google Cloud Workflows), supporting elastic scaling.
- CrewAI-based Orchestration: A multi-agent coordination framework structuring roles and negotiation (e.g., liquidity proposing transfers, FX raising alerts, compliance validating rules).
- Enterprise Middleware Integration: Agents embedded into existing treasury management systems (TMS) or ERP platforms, with the orchestrator as a rules engine (e.g., Drools, Camunda).
- Hybrid Model: Forecasting and liquidity agents in the cloud for scalability, with compliance-sensitive orchestration on-premises for regulatory control.

These alternatives demonstrate that the framework is both model-agnostic and technology-neutral, allowing organizations to adapt implementation to their operational, security, and compliance needs.

3.3.3. Leveraging Pre-Trained Models and Agentic Protocols

Recent advances in agentic AI suggest opportunities to accelerate implementation using pre-trained models and standardized protocols.

- Pre-trained models reduce development costs and enable rapid deployment. For example, LLMs fine-tuned on financial corpora can be repurposed for sentiment analysis or compliance checking, while open-source forecasting models can support predictions of liquidity and yield.
- Model Context Protocol (MCP) provides a standardized way for AI agents to access external data sources, tools, and workflows securely (Hou *et al.*, 2025). MCP ensures consistent context sharing and reduces integration complexity across heterogeneous systems.
- Governance protocols such as MI9 (C. L. Wang *et al.*, 2025) and SAGA (Syros *et al.*, 2025) introduce runtime monitoring, identity management, and risk scoring for agent systems, making them more suitable for regulated domains like treasury.

Table 2 summarizes the key advantages and disadvantages of using pretrained models and standardized agent protocols in enterprise treasury applications.

Table 2. Pre-trained Models and Agentic Protocols: Pros and Cons

Aspect	Advantages	Challenges
Pre-trained Models	Faster development; leverage large corpora; reduced training costs; immediate access to state-of-the-art NLP and forecasting	Domain mismatch; need fine-tuning; explainability and auditability concerns; possible bias in training data
Model Context Protocol (MCP)	Standardized access to external tools/data; consistent context sharing; reduces integration complexity	Emerging standard; limited adoption; requires adaptation for financial APIs and regulatory datasets.
Governance Protocols (MI9, SAGA, etc.)	Provides identity, runtime governance, risk monitoring, and audit trails; improves trust and compliance.	Increases system complexity, adds performance overhead, and governance frameworks are still evolving.

Source: (Hou *et al.*, 2025); (C. L. Wang *et al.*, 2025); (Syros *et al.*, 2025); (Hernes *et al.*, 2024); (Ayatollahi & Jafari, 2025); Authors' analysis.

Integrating these approaches would allow enterprises to combine robust pre-trained intelligence with compliance-grade oversight, aligning innovation with regulatory and operational requirements.

4. RESULTS AND DISCUSSION

Secondary analysis of The Harmony Gap study confirms the scale of financial inefficiencies in global enterprises. Organizations lose an average of $98.5 million annually due to "disharmony" in their money lifecycle, with costs concentrated in cyberthreats, fraud, and regulatory complexity.

Figure 2 summarizes the distribution of these costs, showing cyberthreats alone averaging $31.7M in annual losses, followed by fraud ($21.6M), regulatory compliance ($14.9M), and operational inefficiencies ($11M). Smaller but non-negligible sources include human errors, payment friction, skills gaps, illiquidity, and reputational losses.

Figure 2. Baseline Findings: Annual Financial Disharmony Costs

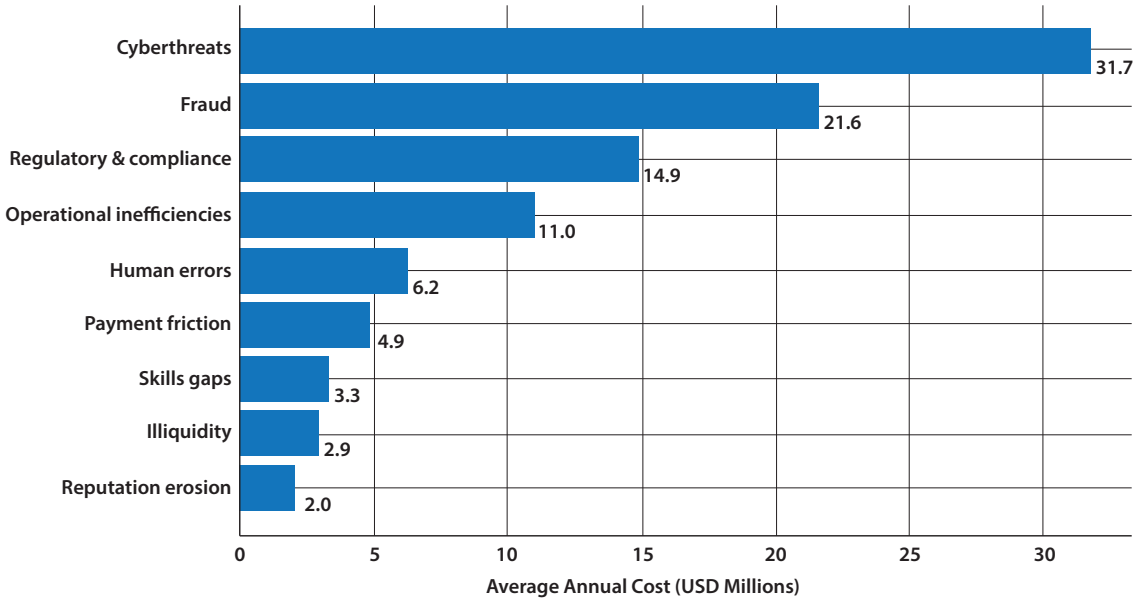

Source: Author-generated, based on data from the Harmony Gap report, FIS & Oxford Economics, 2025

Notably, 51% of surveyed organizations identified the money-in-motion phase (payments, transfers, settlements) as the most vulnerable stage of the lifecycle. Industry differences are significant:

- Tech/Big Tech firms are disproportionately impacted by fraud (28%).
- Insurance companies face heightened compliance burdens (20%).
- Fintechs report greater cyber-risk exposure (35%).
- Building on Harmony Gap evidence, our secondary analysis highlights problem areas particularly relevant to treasury and digital asset operations. Figure 3 depicts these areas.
- Idle liquidity: $5B–$10B in daily idle cash across large treasuries, leading to ~$30M annual lost yield.
- FX risk exposure: up to 80% of flows unhedged due to manual or static processes.
- Operational lags: 3–5 day delays in reallocating funds across entities, creating inefficiencies and opportunity costs.
- Digital asset inefficiencies: $2B+ idle capital trapped in tokenized cash and wallets, requiring 5× overfunding due to lack of real-time mobility.

Figure 3. Identified Problem Areas in Treasury and Digital Assets

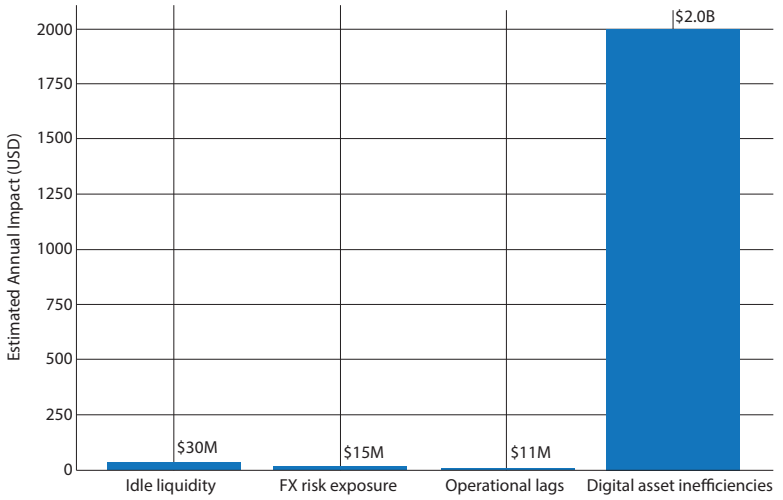

Source: Author-generated, partially based on Harmony Gap data and authors' analysis

These estimates are derived from secondary analysis and assumptions, including estimated yield losses, extrapolated FX-related costs, reported operational inefficiencies, and calculated overfunding needs for digital asset immobility. These insights correspond with Harmony Gap findings that operational inefficiency, compliance issues, and money-in-motion friction are the primary sources of disharmony, strengthening the case for adaptive AI-driven solutions. Table 3 benchmarks the projected benefits of the proposed multi-agent framework against the inefficiencies of the Harmony Gap.

Table 3. Comparative Benefits of Multi-Agent AI Framework

Challenge (Harmony Gap / Treasury)	Cost/Impact Today	Multi-Agent AI Benefit
Idle liquidity & underperforming deposits	$30M annual yield lost; $5–10B idle daily balances	Dynamic reallocation to MMFs/T-bills; optimized yield and liquidity
FX risk exposure	80% of flows unhedged	FX Agent dynamically hedges using real-time curves and volatility thresholds
Money-in-motion friction	51% of firms cite as the top source	Automated settlement, real-time liquidity routing, faster cash mobility
Operational inefficiencies	$11M annually per firm	Forecasting + automation eliminates delays and manual errors
Compliance/reporting burden	$14.9M annually	Embedded auditability: automated reporting that complies with disclosure rules
Digital asset inefficiency	$2B+ idle trapped liquidity	Unified ledger coordination; real-time reserve deployment

Source: Authors' framework design; Harmony Gap Report, 2025

The Harmony Gap study highlights a baseline inefficiency of nearly $100M per enterprise annually, with the bulk of losses linked to cyber, fraud, compliance, and operational inefficiencies. Secondary analysis narrows this further to treasury-specific pain points, including idle liquidity, unhedged FX exposure, and inefficiencies in both traditional and digital asset flows.

The integration of Harmony Gap insights with multi-agent AI design highlights a clear path from problem to solution. The Harmony Gap measures the financial scale of inefficiencies, while our framework suggests agentic intelligence as a way to restore agility in treasury and digital assets.

From a business agility perspective, the system enables organizations to:

- Respond dynamically to market volatility through AI-based hedging and liquidity routing.
- Reduce reliance on manual processes, thereby lowering human error and operational risk.
- Meet regulatory expectations proactively with built-in compliance reporting.
- Enhance customer experience with instant, traceable payments and disbursements.

Liquidity and FX agents mitigate idle balances and unhedged exposure, while the orchestrator ensures compliance and auditability. Importantly, the framework is adaptable: enterprises may implement it through microservices, cloud-native orchestration (Azure Durable Functions, AWS Step Functions), or CrewAI-based multi-agent coordination. However, risks remain in implementation, particularly around cybersecurity, integration costs, and talent gaps (challenges also noted in Harmony Gap). Success depends on phased deployment, robust governance, and alignment with regulatory standards.

While this study introduces a conceptual multi-agent AI framework for adaptive treasury and capital optimization, several limitations should be noted. First, the analysis relies on secondary data, which restricts empirical validation of the proposed design. Future research should test the framework through simulations, prototypes, or pilot implementations. Second, the effectiveness of the agents depends heavily on the quality, availability, and timeliness of financial and market data. Incomplete or inconsistent datasets may reduce accuracy and reliability. Third, integrating with existing treasury systems poses significant challenges, including interoperability with legacy infrastructure, compliance with regulatory requirements, and cybersecurity risks. Finally, over-reliance on automation may create operational vulnerabilities if human oversight is reduced. These limitations highlight the need for careful governance and phased adoption when applying multi-agent AI in real-world treasury operations.

5. CONCLUSION

Modern enterprises face increasing pressures from unpredictable markets, intricate regulations, and disjointed financial systems (Yulfajar *et al.*, 2025). As the Harmony Gap study highlights, inefficiencies in the money lifecycle cost organizations nearly $100 million annually, with the most significant delays occurring in money-in-motion processes. These insights emphasize the urgent need for solutions that improve resilience, efficiency, and agility in capital management.

This paper presents a multi-agent AI framework for adaptive treasury and capital optimization in response to these challenges. By integrating specialized agents for liquidity, FX, interest rates, and market sentiment within a rule-based recommendation system, the framework seeks to reduce idle balances, dynamically hedge exposures, lower transaction costs, and ensure real-time compliance. Comparative analysis shows that this approach can provide significant operational and financial benefits.

Adopting this model directly enhances business agility: companies that can reallocate capital quickly, adjust to market changes, and maintain transparent compliance will be best positioned to stay competitive. Future research should empirically test the framework through simulations and pilot projects, with a focus on cybersecurity, data governance, and integration costs. By connecting industry-documented inefficiencies with emerging AI capabilities, this study presents multi-agent intelligence to create more agile, resilient, and innovative treasury and capital management.

LITERATURE

Atienza-Barba, M., del Río, M. de la C., Meseguer-Martínez, Á., & Barba-Sánchez, V. (2024). Artificial intelligence and organizational agility: An analysis of scientific production and future trends. *European Research on Management and Business Economics, 30*(2), 100253.

Ayatollahi, M., & Jafari, S. M. (2025). Orchestrated Intelligence: *A Multi-Agent Architecture for High-Frequency Quantitative Risk Management and Execution in Complex Options Markets*.

Dodnikov, R. (2024). Workflow Optimization in Finance Using Multi-Agent Systems. *Innovative Information Technologies for Economy Digitalization (IITED), 1*, 182–188.

Donepudi, P. K., Banu, M. H., Khan, W., Neogy, T. K., Asadullah, A. B. M., & Ahmed, A. A. A. (2020). Artificial intelligence and machine learning in treasury management: a systematic literature review. *International Journal of Management, 11*(11).

FIS, & Oxford Economics. (2025). *The Harmony Gap: Finding the Financial Upside in Uncertainty*. https://fisglobal.com/the-harmony-gap

Fosso Wamba, S., Queiroz, M. M., Guthrie, C., & Braganza, A. (2022). Industry experiences of artificial intelligence (AI): benefits and challenges in operations and supply chain management. In *Production planning & control* (Vol. 33, Issue 16, pp. 1493–1497). Taylor & Francis.

Hernes, M., Korczak, J., Krol, D., Pondel, M., & Becker, J. (2024). Multi-agent platform to support trading decisions in the FOREX market. *Applied Intelligence, 54*(22), 11690–11708.

Hou, X., Zhao, Y., Wang, S., & Wang, H. (2025). Model Context Protocol (MCP): Landscape, Security Threats, and Future Research Directions. *ArXiv E-Prints*, arXiv–2503.

Johnston, M. P. (2014). Secondary data analysis: A method of which the time has come. *Qualitative and Quantitative Methods in Libraries, 3*(3), 619–626.

Kocot, M. (2023). Organizational Agility in the Aspect of Sustainable Development of the Organization. *Scientific Papers of Silesian University of Technology. Organization & Management/Zeszyty Naukowe Politechniki Slaskiej. Seria Organizacji i Zarzadzanie, 187*.

Leitner, G., Singh, J., van der Kraaij, A., & Zsámboki, B. (2024). The rise of artificial intelligence: benefits and risks for financial stability. *Financial Stability Review, 1*.

Maple, C., Szpruch, L., Epiphaniou, G., Staykova, K., Singh, S., Penwarden, W., Wen, Y., Wang, Z., Hariharan, J., & Avramovic, P. (2023). The AI revolution: Opportunities and challenges for the finance sector. *ArXiv Preprint ArXiv:2308.16538*.

Rastogi, S., & Pandita, D. (2025). Driving entrepreneurial success: navigating AI-driven transformation through workforce agility and sustainability. *Journal of Innovation and Entrepreneurship, 14*(1), 75.

Shirzad, A., & Rahmani, A. (2024). Smart Treasury: Leveraging Artificial Intelligence and Robotic Process Automation for Financial Excellence. *Knowledge Economy Studies, 1*(1), 65–86.

Syros, G., Suri, A., Nita-Rotaru, C., & Oprea, A. (2025). SAGA: A Security Architecture for Governing AI Agentic Systems. *ArXiv E-Prints*, arXiv–2504.

Wang, C. L., Singhal, T., Kelkar, A., & Tuo, J. (2025). MI9–Agent Intelligence Protocol: Runtime Governance for Agentic AI Systems. *ArXiv E-Prints*, arXiv–2508.

Wang, X., Lin, X., & Shao, B. (2022). How does artificial intelligence create business agility? Evidence from chatbots. *International Journal of Information Management, 66*, 102535.

Yulfajar, A., Noor, G. M., & Putranto, R. S. (2025). The Impact of Financial Market Instability on Economic Growth and Long-Term Investment. *Advances in Economics & Financial Studies, 3*(1), 43–55.