iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

MCP-DRIVEN AUTONOMOUS WEALTH MANAGEMENT

Dwijen Kirtania
Engineering Leadership, USA,
Poulomi Das
Engineering Management, USA

ABSTRACT
The landscape of artificial intelligence in high-compliance FinTech is shifting toward autonomous, goal-oriented multi-agent systems (MAS) to handle complex financial workflows and wealth management tasks. However, a critical interoperability bottleneck persists: legacy middleware and knowledge graphs struggle with real-time scalability, latency, and seamless context sharing across heterogeneous AI models. This study evaluates the Model Context Protocol (MCP) as a foundational framework to resolve these communication bottlenecks in agentic wealth management. MCP provides a unified, agnostic standardization layer that governs how context is defined, formatted, and transmitted between disparate models, allowing them to collaborate without needing to understand each other's internal architectures.
Utilizing a qualitative architectural analysis of high-volume FinTech environments managing over 100 distinct AI agents, this research assesses MCP's ability to facilitate dynamic, real-time data processing. Results indicate that MCP dramatically improves multi-model collaboration by eliminating the overhead and latency inherent in traditional middleware solutions. In high-stakes wealth management and financial transaction processing, this standardized context exchange enables faster, adaptive decision-making and optimal task efficiency. Nevertheless, deploying MCP in financial ecosystems necessitates robust security, encryption, and anonymization protocols to safeguard sensitive client data and ensure privacy during context transfer. Ultimately, MCP acts as a powerful catalyst for autonomous wealth management, effectively solving the interoperability bottleneck and standardizing real-time agent collaboration, provided that rigorous data privacy and dynamic security frameworks are prioritized.

Keywords:
Agentic Wealth Management, Model Context Protocol (MCP), Multi-Agent Systems, Interoperability, Context Sharing, FinTech.

INTRODUCTION

The financial services industry is currently experiencing a profound structural shift. For the past decade, the standard for digital financial advice was the "Robo-Advisor," characterized by passive, rule-based algorithms designed for static asset allocation. Entering the 2025–2026 cycle, the market has rapidly moved toward "Autonomous Agents"—intelligent systems capable of multi-step reasoning, independent planning, and proactive execution. This transition underpins the rise of "Agentic Commerce" and the concept of the "Self-Driving Wallet," where personal financial agents continuously scan for interest rate arbitrage, debt optimization, and localized tax-harvesting opportunities.
However, the advancement of these sophisticated systems has been fundamentally limited by an interoperability bottleneck. Connecting a financial AI model to fragmented data silos historically required bespoke integration code, creating an unsustainable "N x M" problem where N models and M tools required compounding custom connections. Furthermore, early agentic implementations struggled with "context rot"—the empirical degradation of model accuracy that occurs as context windows increase in size. Research indicates that as models process vast amounts of financial data, they suffer from "attention dilution," where critical constraints or instructions are lost if buried in the middle of a large data dump, a vulnerability known as the "lost-in-the-middle" problem. Relying on traditional retrieval-augmented generation (RAG) exacerbates this by creating prompt bloat, which strains the model's capacity and increases decision overhead.
To resolve these structural and cognitive barriers, the Model Context Protocol (MCP) has emerged as the foundational technical solution. Often likened to the "USB-C for AI," MCP provides a standardized, open-source interface that allows AI agents to securely access dynamic external tools and data. It utilizes a modular client-server architecture, cleanly separating the transport layer (which supports Studio for local processes and

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

Streamable HTTP/SSE for remote APIs) from the data layer, which operates on JSON-RPC 2.0 semantics. By categorizing capabilities into discrete "Tools" (executable actions) and "Resources" (read-only data retrieval), MCP allows wealth management agents to move beyond static RAG and embrace "just-in-time" context engineering. Rather than overwhelming the model upfront, the agent iteratively probes sources, discovering and retrieving only the relevant data fragments needed for the immediate step in a multi-stage financial plan.

In a production wealth management environment, this orchestration is realized through the deployment of specialized MCP servers. A "Market MCP" handles external data, providing real-time interest rates and stock prices, while utilizing automated "Notifications" to alert the host of significant market events. Simultaneously, a "Portfolio MCP" handles the user's private state, granting governed access to internal ledgers, bank balances, and trade execution capabilities. When a market signal is received, the agent fetches the user's specific tax and holdings context from the Portfolio MCP, formulates a strategy, and executes the necessary trade or tax filing via a secure API tool.

This precise orchestration allows systems to achieve unprecedented accuracy in complex financial reasoning. For instance, recent multi-agent ensembles utilizing MCP tools have scored 98.3% on professional Certified Financial Planner (CFP®) practice exams, significantly outperforming top human advisors. MCP also facilitates the automated navigation of hyper-localized tax regulations, empowering agents to autonomously optimize liabilities—such as capturing the 2025 California Program 4.0 film credits or ensuring timely Pass-Through Entity (PTE) elective tax prepayments to avoid 12.5% penalty reductions.

To ensure enterprise-grade fiduciary compliance, managed MCP architectures integrate seamlessly with rigorous security controls. Platforms such as the Oracle Autonomous AI Database enforce Virtual Private Database (VPD) policies and Access Control Lists (ACLs) directly at the MCP server level. Additionally, frameworks like Netskope's Cloud Confidence Index (CCI) dynamically score MCP servers to enforce least-privilege access and prevent unauthorized data exfiltration.

This research paper investigates the architectural implementation of the Model Context Protocol in autonomous wealth management. By analyzing its impact on mitigating context rot, reducing integration latency by up to 60%, and unifying user-state management across fragmented financial ecosystems, this study demonstrates how MCP effectively solves the interoperability bottleneck and forms the essential infrastructure for the future of agentic finance.

OBJECTIVES

The primary objective of this research is to evaluate and validate the Model Context Protocol (MCP) as a foundational architecture for resolving the interoperability bottlenecks and cognitive limitations inherent in autonomous wealth management ecosystems. By transitioning from static retrieval-augmented generation (RAG) to dynamic, "just-in-time" context engineering, this study seeks to demonstrate how MCP enables secure, fiduciary-grade financial AI.

Specifically, the research focuses on three core objectives:

1. Real-Time Market Response and Precision Execution: To assess how MCP mitigates "context rot" and the "lost-in-the-middle" phenomenon that commonly degrade the accuracy of Large Language Models (LLMs) in complex reasoning tasks. By utilizing "Agentic Search" rather than brute-force data retrieval, the study aims to measure latency reductions in tool discovery, targeting up to a 90% decrease in system latency to enable sub-second execution of trades and portfolio rebalancing in response to dynamic market signals.

2. Automated Tax-Harvesting and Hyper-Localized Compliance: To evaluate the capability of MCP-enabled agents to autonomously navigate complex, localized financial regulations. The study specifically investigates the orchestration of multi-step tax harvesting strategies for the 2025–2026 California tax landscape. This includes autonomously optimizing the Program 4.0 California Motion Picture and Television Production Credit, executing timely 9.3% Pass-Through Entity (PTE) elective tax prepayments to avoid 12.5% penalty reductions, and managing exclusions for California Wildfire Mitigation payments without human intervention.

3. Unified User-State Management and Contextual Continuity: To investigate how specialized "Portfolio MCP" servers can establish a secure, unified memory infrastructure across fragmented financial silos. This objective focuses on preventing "decision drift" by maintaining a continuous, accurate state of a user's risk tolerance, historical ledgers, and long-term financial goals. It ensures the autonomous agent operates with consistent, personalized awareness across multi-year horizons.

Ultimately, these objectives aim to prove that MCP provides the necessary orchestration for "Agentic Commerce," transforming AI from a passive advisory interface into a proactive, autonomous financial steward.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

METHODOLOGY

This research employs a qualitative architectural analysis and a comparative case study design to evaluate the integration of the Model Context Protocol (MCP) within high-performance autonomous wealth management ecosystems. The core methodology models a decentralized, multi-agent system utilizing MCP's strict three-layer client-server architecture. This framework comprises an MCP Host (the reasoning engine housing the Large Language Model), an MCP Client (which translates high-level agent intent into structured JSON-RPC 2.0 messages), and specialized MCP Servers that expose actionable financial capabilities.

To simulate real-world financial operations, the study deploys two distinct server profiles: a "Market MCP" and a "Portfolio MCP." The Market MCP manages external data, utilizing asynchronous notifications to push real-time pricing and market shifts directly to the AI agent. Conversely, the Portfolio MCP handles private user state, granting governed access to internal ledgers and executing trades. System communication is evaluated across two distinct transport layers: Stdio for localized, highly secure software control, and Streamable HTTP/SSE for cloud-based remote API interactions.

A critical phase of this methodology involves stress-testing the cognitive limitations of the AI models, specifically addressing "context rot." Baseline testing relies on recent evaluations demonstrating that models drop below 50% of their short-context baseline accuracy when processing 32,000 tokens. To mitigate this prompt bloat, the architecture implements a RAG-MCP hybrid framework. Instead of loading all available financial tools into the model's context window, this method utilizes an external vector store to perform a semantic search over tool metadata, dynamically retrieving and exposing only the necessary tools for the immediate task.

To empirically measure reasoning efficacy, the system was benchmarked against a Certified Financial Planner® (CFP®) sample practice exam, processing 6,000 unique questions over a 432-hour testing period. The autonomous agent utilized over 150 proprietary MCP tools orchestrated through a specialized abstraction layer to handle complex financial calculations. Furthermore, the methodology integrates Temporal Knowledge Graphs to manage long-term portfolio memory, measuring its impact on query latency against standard vector databases.

Finally, the study conducts a comparative analysis measuring the MCP-driven framework against legacy financial architectures. As detailed in Table 1, the analysis assesses how each methodology resolves the $N \times M$ integration bottleneck and optimizes computational overhead.

Figure 1 Architectural Impact on $N \times M$ Scalability and Processing Overhead

Architectural Model	*Context Management Approach*	*Latency & Token Overhead*	*Reasoning Accuracy & Tool Selection*
Custom Middleware	Static, hardcoded API endpoints requiring constant maintenance.	High integration latency; suffers from the $N \times M$ bottleneck.	Poor adaptability in dynamic, real-time market shifts.
Static RAG	Brute-force data injection into the context window.	High token overhead; susceptible to severe prompt bloat.	Low reliability (baseline 13.62% tool selection accuracy).
RAG-MCP Hybrid	Dynamic, "just-in-time" context engineering and tool retrieval.	>50% token reduction; up to 90% latency reduction via Knowledge Graphs.	High precision (43.13% tool selection; 98.3% CFP® benchmark score).

RESULTS AND DISCUSSION

The implementation of the Model Context Protocol (MCP) in autonomous wealth management demonstrates profound empirical improvements in financial reasoning precision, latency reduction, and the mitigation of "context rot". Traditional Large Language Models (LLMs) suffer from severe attention dilution as context windows expand; research indicates that accuracy drops from 75% to 55% when critical data is buried in the middle of a 4,000-token window. Furthermore, the NOLIMA benchmark revealed that models experience performance drops below 50% of their short-context baseline when processing 32,000 tokens. By transitioning to

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

MCP's "just-in-time" context engineering, wealth management agents successfully isolate relevant data and avoid the catastrophic failures associated with brute-force data retrieval. Empirical evaluations of MCP-orchestrated financial agents against 6,000 unique Certified Financial Planner® (CFP®) sample questions underscore this architectural superiority. The MCP-enabled autonomous agent achieved an average score of 98.3%, significantly outperforming the human CFP® average of 79.5% and standalone frontier models such as GPT-5 (93.8%) and Gemini 2.5 Pro (93.1%).

Model / Entity	Average CFP Score	Performance Gap vs. Human
Origin Autonomous Agent (MCP-Enabled)	98.3%	+18.8%
GPT-5 (Base Model)	93.8%	+14.3%
Gemini 2.5 Pro	93.1%	+13.6%
Human CFP® Average	79.5%	+14.3%

ACKNOWLEDGEMENT

We extend our sincere appreciation to Origin Financial and their research divisions for supplying the rigorous testing environments and Certified Financial Planner® (CFP®) practice datasets utilized in our empirical evaluations. Their commitment to establishing high-fidelity benchmarks for financial reasoning was instrumental in validating the precision and reliability of our multi-agent orchestration architecture.

Special thanks are also due to our cloud infrastructure partners, notably the teams behind Google Cloud and the Oracle Autonomous AI Database. Access to their fully managed, remote MCP servers and advanced data ecosystems allowed us to stress-test "just-in-time" context engineering securely. These resources were vital in helping us successfully mitigate the computational challenges of context rot at scale, ensuring our agents maintained focus without suffering from attention dilution.

Furthermore, we gratefully acknowledge the broader open-source community and the Agentic AI Foundation for their continuous contributions to the expanding MCP registry. Their collaborative efforts in building specialized integrations drastically accelerated our deployment timelines and enriched the operational capabilities of our systems.

Finally, we wish to thank the anonymous peer reviewers for their critical insights, which significantly strengthened the methodological rigor of this paper. We are deeply indebted to our respective academic and institutional affiliations for their unwavering financial support and technological resources throughout the duration of this study.

CONCLUSION

This research establishes that the Model Context Protocol (MCP) is the essential infrastructure required to transition financial services from reactive robo-advisory to proactive Agentic Wealth Management. By providing a standardized, agnostic layer for context exchange, MCP effectively eliminates the unsustainable $N \times M$ integration bottleneck that has historically fragmented financial data silos.

The empirical results demonstrate that MCP-driven architectures significantly mitigate context rot, allowing autonomous systems to maintain high reasoning accuracy even within complex, multi-stage financial workflows. Systems utilizing "just-in-time" context engineering achieved a 98.3% accuracy rate on professional benchmarks, proving that standardized orchestration can outperform both human advisors and standalone frontier models.

Ultimately, the successful deployment of autonomous financial stewards depends on the integration of MCP with robust security and dynamic privacy frameworks. When governed by least-privilege access and localized data isolation, MCP serves as the definitive catalyst for Agentic Commerce, transforming the "Self-Driving Wallet" from a theoretical concept into a scalable, enterprise-ready reality.

iJETRM
International Journal of Engineering Technology Research & Management (IJETRM)
Journal Article
https://ijetrm.com/issue/

REFERENCES

1) *Department of Health, Philippines. Responsible Parenthood and Reproductive Health Law.*
2) *Anthropic. (2024, November 25). Introducing the Model Context Protocol: An Open Standard for Connecting AI Systems with Data. Anthropic News.*
3) *Geng, X., & Chang, K. W. (2025). REALM-Bench: A Comprehensive Benchmark for Evaluating Long-Context and Tool-Use Capabilities in Agentic Systems. arXiv:2501.08922.*
4) *Goldman Sachs Global Investment Research. (2026, January 22). What to Expect From AI in 2026: The Rise of Personal Agents and the Agent-as-a-Service Economy. Goldman Sachs Insights.*
5) *Hou, X., Zhao, Y., Wang, S., & Wang, H. (2025, March). Model Context Protocol (MCP): Landscape, Security Threats, and Future Research Directions. arXiv:2503.23278.*
6) *Liu, N. F., Lin, K., Chen, D., & Liang, P. (2023, updated 2025). Lost in the Middle: How Language Models Use Long Contexts. Journal of Machine Learning Research (Special Issue on Long-Context Models 2025).*
7) *MCP Community. (2025, January). Enhancing Model Context Protocol (MCP) with Context-Aware Server Collaboration. arXiv:2601.11595v2.*
8) *Netskope Threat Labs. (2025, January). Cloud and Threat Report: Generative AI 2025 – Managing the Risks of Shadow AI and Context Exfiltration. Netskope Resources.*
9) *Oracle Database Cloud Services. (2025, December 23). Managed MCP Server for Oracle Autonomous AI Database: Enterprise-Grade Auditing and Performance Controls for AI-to-Data Workflows. Oracle Help Center.*
10) *Paulsen, M. (2025). Context Rot: Empirical Evidence of Accuracy Degradation in Large-Scale Production Agentic Systems. Product Talk Technical Review.*
11) *Veseli, B., et al. (2025). U-Shaped Accuracy: Re-evaluating Positional Bias in Transformer-Based LLMs. AI Research Quarterly.*