International Journal of Innovative Research in Engineering and Management (IJIREM)
ISSN (Online): 2350-0557, Volume 13, Issue 4, August 2026
DOI: https://doi.org/10.55524/ijirem.2026.13.4.5
Article ID IJIRE-1468, Pages 35-42
www.ijirem.org

Model Context Protocol (MCP) or Agent-to-Agent (A2A) — Defining Right Communication Protocol for Agentic AI in Financial Services Use Cases

(A Comparative Analysis of Agent Interoperability Standards for Regulated Financial Services Institutions)

Aditya Ghosh[1], and Siddhartha Kumar Ghosh[2]

[1] Student, Narayana eTechno School, Kalyan, Maharashtra, India
[2] Cluster Leader, Financial Services Sector, IBM, India/South Asia

Correspondence should be addressed to Aditya Ghosh; skg1609@gmail.com

Received: 14 July 2026; Revised: 28 July 2026; Accepted: 10 August 2026

ABSTRACT- The rapid adoption of large language model (LLM)-based agentic systems in financial services has created an urgent architectural question for enterprise technology leaders: should an institution standardize on the Model Context Protocol (MCP), introduced by Anthropic in November 2024, or the Agent2Agent (A2A) protocol, introduced by Google in April 2025, when designing agentic AI infrastructure? This paper undertakes a structured comparative analysis of both protocols through the lens of financial-services use cases such as fraud detection, credit decisioning, model risk management, regulatory compliance, treasury operations, and agentic commerce and payments. Drawing on protocol specifications, peer-reviewed and preprint literature, and documented industry deployments at institutions including JPMorgan Chase, Goldman Sachs, Commonwealth Bank, Visa, and Mastercard, the study finds that MCP and A2A occupy complementary layers of the agentic AI stack rather than competing for the same role: MCP is a vertical protocol governing how an individual agent accesses tools, data, and enterprise systems, while A2A is a horizontal protocol governing how autonomous agents owned by different teams, vendors, or institutions discover, negotiate with, and delegate tasks to one another. To move beyond a purely narrative comparison, the paper adds a literature-supported case-study table of documented financial-services deployments and a rubric-based comparative scoring exercise across six operational criteria. The results support a decision framework in which mature agentic AI architectures in regulated finance converge toward a layered deployment: MCP grounds individual agents in trusted enterprise data, and A2A orchestrates cross-institutional and cross-vendor collaboration, with emerging payment-layer extensions such as the Agent Payments Protocol (AP2) and x402 bridging agent coordination to settlement. The paper closes by identifying research gaps around protocol security maturity, regulatory liability for autonomous agent transactions, and the absence of controlled empirical benchmarking in the financial-services domain.

KEYWORDS- Model Context Protocol; Agent2Agent Protocol; Agentic AI; financial services; multi-agent systems; agent interoperability; agentic payments; AI governance; agent security.

I. INTODUCTION

Financial institutions have moved rapidly from experimenting with generative AI chatbots to deploying autonomous, goal-directed agentic systems that plan, invoke tools, and act with limited human supervision across multi-step workflows [1][2]. Industry reporting indicates that JPMorgan Chase operates several hundred production AI use cases on an enterprise platform touching trillions of dollars in daily transaction volume, that Goldman Sachs pairs thousands of engineers with coding and operational agents, and that banks such as Commonwealth Bank, HSBC, Citi, UBS, DBS, and ING have reported measurable reductions in fraud losses and operating costs from agentic deployments [3]. Reporting on banking-sector security posture further notes that AI-agent adoption in banks is outpacing the maturity of the security controls surrounding it [4]. As these systems scale beyond single-agent copilots into networks of specialized agents — a KYC agent, a fraud-detection agent, a treasury agent, a compliance-documentation agent — the question of how such agents communicate with tools, data, and each other has become a first-order architectural decision.

Two open protocols have emerged as the leading candidates for standardizing this communication. The Model Context Protocol (MCP), released by Anthropic in late 2024, defines a client-server interface through which an agent discovers and invokes external tools, retrieves contextual resources, and uses templated prompts [1]. The Agent2Agent protocol (A2A), released by Google in April 2025 with the backing of more than fifty technology and consulting partners, defines a peer-to-peer standard through which independently developed agents communicate their capabilities via Agent Cards and negotiate stateful tasks with one another [2]. Because both protocols are frequently discussed together in vendor literature [5-10], and because financial-services

architecture teams are under pressure to commit to a standard ahead of regulatory and audit cycles, practitioners often frame the decision as an either/or choice: MCP or A2A.

This paper argues that framing is not a right depiction to pick an either/or choice. It undertakes a structured literature-based comparison of the two protocols, maps each to representative financial-services use cases, adds a literature-informed case-study and rubric-based empirical assessment, and synthesizes the emerging consensus that MCP and A2A solve problems at different layers of the same stack and are typically deployed together rather than as substitutes.

A. *Objectives*

To characterize the architectural, security, and governance differences between MCP and A2A based on current literature and protocol documentation.

To map each protocol to concrete financial-services agentic AI use cases where it offers clearer fit.

To assess documented financial-services deployments through a structured case-study and rubric-based comparative exercise.

To propose a decision framework that financial-services technology leaders can apply when architecting agentic AI systems.

B. *Structure of the paper*

Section 2 reviews the academic and industry literature on agent interoperability protocols and their application to financial services. Section 3 describes the research methodology. Section 4 presents the comparative architectural results and their implications for financial-services architecture, security, and regulation. Section 5 presents an empirical case-study and rubric-based comparative assessment. Section 6 concludes the paper, and Section 7 identifies directions for future research.

II. LITERATURE REVIEW

A. *Classical foundations of agent communication*

Standardized agent communication predates the current LLM-based agent wave by three decades. The Knowledge Query and Manipulation Language (KQML), developed under DARPA's Knowledge Sharing Effort, defined an extensible set of message performatives allowing agents to exchange knowledge independent of a specific ontology or transport. The Foundation for Intelligent Physical Agents (FIPA) subsequently defined an Agent Communication Language (ACL) grounded in speech-act theory, together with a library of interaction protocols for contract-net bidding, auctions, and query-reference exchanges [11,12]. Formal analyses of FIPA-ACL's semantics clarified how its performatives should be interpreted against the mental states of communicating agents [13]. Wooldridge and Jennings' foundational characterization of intelligent agents — autonomy, social ability, reactivity, and pro-activeness — remains a reference point for evaluating modern agent architectures, including LLM-based ones [14]. Despite their conceptual rigor, neither KQML nor FIPA-ACL achieved broad adoption in today's LLM-based agent ecosystems, a gap that MCP and A2A were designed to fill using web-native transports and schema-based tool discovery rather than formal speech-act semantics.

B. *The emergence of MCP and A2A*

Anthropic introduced MCP as an open standard intended to replace bespoke, point-to-point tool integration with a single, reusable interface between AI applications and external tools, resources, and prompt templates [1]. Within months, the protocol was adopted by OpenAI, Google DeepMind, and Microsoft/GitHub, and industry trackers report thousands of public MCP servers and millions of weekly package installations [15,16]. In December 2025, Anthropic transferred MCP governance to the newly formed Agentic AI Foundation under the Linux Foundation, with platinum members including AWS, Google, Microsoft, Bloomberg, Cloudflare, and OpenAI [16], signaling the protocol's transition from a single-vendor specification to community-governed infrastructure.

Google's A2A protocol emerged several months later to address a different limitation: MCP standardizes how a single agent reaches tools and data, but does not natively address how two independently built, potentially opaque agents — built by different vendors, teams, or competing institutions — discover one another's capabilities and coordinate multi-step work [2]. A2A addresses this through Agent Cards published at a well-known endpoint, together with a stateful Task construct that tracks a unit of delegated work through its lifecycle [8][9].

C. *Comparative and survey literature*

A growing body of comparative literature clarifies the relationship between the two protocols. Several practitioner and preprint sources converge on describing MCP as a vertical protocol connecting a model to tools and data, and A2A as a horizontal protocol connecting agents to other agents, explicitly framing the two as complementary building blocks [5-9]. Broader academic surveys extend the comparison to a wider protocol family, including the Agent Communication Protocol (ACP) and the Agent Network Protocol (ANP), and catalogue distinct architectural approaches: MCP as a JSON-RPC client-server interface for tool invocation, A2A as a peer-to-peer standard for capability-based task outsourcing, and ACP as a REST-native protocol for multimodal agent responses [17]. Related survey work situates MCP and A2A explicitly within the historical arc from symbolic, formally-specified protocols such as KQML and FIPA-ACL toward web-native, LLM-oriented interoperability standards [18,19]. A recurring theme is that the two protocols were designed to interlock rather than compete: production multi-agent systems are reported to use MCP for data and tool access and A2A for cross-agent coordination within the same deployment [7].

D. *Agentic AI in financial services*

A parallel and rapidly growing literature examines agentic AI applications within financial services specifically. Comprehensive surveys characterize agentic AI in finance as distinguished from both traditional algorithmic trading and earlier generative AI by its goal-oriented autonomy, continuous learning, and multi-agent coordination, while cautioning that these properties introduce novel risks around market stability, interpretability, and systemic risk [20]. Complementary work proposes a four-layer architecture for financial agents spanning data perception, reasoning, strategy generation, and execution with human

control [21]. Domain-specific applied studies illustrate concrete multi-agent designs already piloted in financial institutions, including a multi-agent crew architecture for model risk management comprising a judge agent alongside specialized compliance, replication, and outcome-analysis agents, evaluated across open- and closed-source LLMs [22].

Peer-reviewed engineering literature on multi-agent systems in finance predates the current agentic-AI wave and offers methodologically rigorous precedent, even where it does not explicitly discuss MCP or A2A. A hierarchical multi-agent deep reinforcement learning framework for algorithmic trading, published in Expert Systems with Applications, demonstrated that cooperative multi-agent architectures operating across multiple timeframes outperformed single independent agents and benchmark strategies on foreign-exchange data [23]. A related framework combining multi-agent reinforcement learning with temporal feature extraction for trading strategy optimization was likewise published in Expert Systems with Applications [24]. Work presented at the ACM SIGKDD Conference on Knowledge Discovery and Data Mining demonstrated intention-aware communication among cooperating agents for multi-order execution in finance [25]. This body of peer-reviewed, empirically validated multi-agent finance research establishes that coordinated multi-agent designs already deliver measurable performance gains in financial applications, independent of which interoperability protocol is used to implement the coordination — reinforcing the argument that protocol choice is an infrastructure decision layered on top of, rather than a substitute for, sound multi-agent system design.

E. Agentic payments as an emerging extension layer

A distinct but closely related literature addresses agent-initiated payments. Google's Agent Payments Protocol (AP2) extends A2A with cryptographically signed mandates — Intent Mandates and Cart Mandates — that create a non-repudiable audit trail authorizing an agent to transact on a user's behalf and has been paired with Coinbase's x402 protocol to enable stablecoin-denominated micropayments between agents [26, 27]. Card networks have introduced parallel, network-layer mechanisms: Mastercard's Agent Pay issues tokenized Agentic Tokens scoped to a specific agent and merchant policy, while Visa's Trusted Agent Protocol issues a Verified Agent ID together with a signed consumer consent record [28, 29]. Commentators consistently flag that liability frameworks for erroneous or fraudulent agent-initiated transactions remain unresolved, adding the AI platform itself as a new party to a dispute-resolution landscape historically limited to consumers, merchants, and banks [29, 30].

F. Security, governance, and regulatory literature

A substantial and rapidly expanding body of security research documents structural risks in MCP as adopted at scale. Industry security reporting in 2026 documents dozens of disclosed CVEs against MCP implementations, widespread absence of built-in authentication and authorization in the base specification, and attack patterns including tool poisoning, rug pulls, tool shadowing, and credential exposure through insecure configuration files [31-36]. A2A's trust model, which depends on the integrity of published Agent Cards and external identity infrastructure, has drawn comparatively less published security-vulnerability research to date, though researchers note that cross-organizational delegation between opaque agents raises open questions about verifying an agent's authority and provenance before granting access to sensitive systems [37]. Federated Learning (FL) is proposed to integrate MCP and A2A for privacy-preserving and intelligent mobile banking personalization. FL allows decentralized model training on client devices, protecting sensitive user data [41].

A separate peer-reviewed literature addresses AI governance and explainability in regulated finance, which bears directly on how either protocol would need to be governed if deployed. Cross-disciplinary surveys published in Finance Research Letters catalogue explainability requirements for financial AI systems, noting persistent tension between model performance and regulatory interpretability obligations such as those under SR 11-7 and equivalent regimes [38]. A systematic review published in Artificial Intelligence Review similarly finds that explainability remains only partially resolved across financial AI applications [39]. Comparative doctrinal analysis published in the Journal of Banking Regulation examines algorithmic-governance regimes across nine jurisdictions, finding wide variation in how prudential supervisors currently address autonomous, self-directed AI systems in banking [40]. None of this governance literature yet addresses agent-to-agent interoperability protocols specifically, which is a further gap this paper's discussion (Section IV.C) begins to address.

G. Research gap

While the general MCP-versus-A2A comparison is well covered in practitioner literature, and while agentic AI in finance and AI governance in banking are each fast-growing academic subfields in their own right, the literature has not yet produced a study that explicitly frames the protocol choice as a financial-services architecture decision, combines it with a structured case-study and rubric-based comparative assessment, and incorporates the regulatory, security, and payments-specific considerations unique to banking, capital markets, and insurance. This paper addresses that gap.

III. RESEARCH METHODOLOGY

This study adopts a qualitative, comparative desk-research design, supplemented by a structured case-study and rubric-based comparative assessment, appropriate for an emerging technology area where the primary evidence base consists of protocol specifications, preprint and early peer-reviewed literature, and documented industry deployments. This was then further connected with the industry case studies for doing back-testing and hypothesis validation.

A. Source Identification

Sources include research and case studies in this area, industry real deployments based on practices, supplemented by earlier foundational agent-communication literature (KQML, FIPA-ACL) for historical context. These were supplemented with official protocol specifications and governance announcements from Anthropic, Google, and the Linux Foundation's

Agentic AI Foundation, peer-reviewed journal articles from Expert Systems with Applications, Finance Research Letters, Artificial Intelligence Review, and the Journal of Banking Regulation, and financial-trade press covering documented enterprise deployments and independent security-research publications reporting protocol-level vulnerabilities.

B. Inclusion Sriteria

Directly addresses the architecture, specification, or governance of MCP, A2A, or closely related interoperability protocols (ACP, ANP, AP2, x402).
Addresses agentic or multi-agent AI applications, or AI governance, within financial services, banking, payments, or capital markets.
Published or substantively updated between November 2024 and July 2026 for protocol-specific sources, with earlier foundational sources admitted for historical grounding.
Documents a real, named deployment, pilot, CVE disclosure, empirical result, or protocol governance event rather than purely speculative commentary.

C. Comparative Analysis Framework

Extracted material was synthesized against a structured comparison framework covering nine dimensions: originating organization and governance model, architectural layer of operation, transport and message format, core protocol primitives, interaction pattern, capability discovery mechanism, representative financial-services use cases, documented security posture, and payment-layer extensibility (Section IV.A).

D. Empirical Case-Study And Rubric Methodology

To move beyond narrative synthesis, two additional structured instruments were constructed. First, a case-study table (Table 2) compiles documented, named financial-services deployments, coded against which protocol layer they principally exercise and the outcome reported in the source. Second, a comparative rubric (Table 3) scores each protocol from 1 (weak) to 5 (strong) across six operational criteria derived from Section IV.A, with each score anchored to a specific, cited finding rather than assigned by unweighted expert judgment alone. This rubric is explicitly a literature-informed qualitative synthesis rather than a controlled, quantitative benchmark; it is reported as such, and its limitations are discussed in Section III.E and Section VII.

E. Limitations

As a desk-research synthesis rather than an empirical or experimental study, this paper does not independently benchmark protocol latency, throughput, or security posture through controlled testing; it relies on the accuracy of the secondary and industry sources reviewed, and the rubric in Table 3 reflects qualitative synthesis rather than measured data. The literature base for both protocols is still maturing, is weighted toward vendor and practitioner commentary rather than peer-reviewed publication and evolves quickly; the source material should be understood as reflecting the state of practice through mid-2026. Section 7 identifies the controlled empirical studies needed to address these limitations.

IV. RESULTS AND DISCUSSION

A. Architectural comparison

Table 1 summarizes the comparative findings across the nine dimensions defined in the methodology.

Table 1: Comparative summary of MCP and A2A:

Dimension	Model Context Protocol (MCP)	Agent2Agent Protocol (A2A)
Originator / governance	Anthropic (Nov. 2024); transferred to the Agentic AI Foundation, Linux Foundation (Dec. 2025)	Google (Apr. 2025), with 50+ enterprise and consulting partners
Layer of operation	Vertical: agent-to-tool/data	Horizontal: agent-to-agent
Transport / format	JSON-RPC 2.0 over studio or SSE	HTTP/S with JSON-LD Agent Cards
Core primitives	Tools, Resources, Prompts	Agent Cards, Tasks, Messages
Interaction pattern	Hierarchical client-to-server	Peer-to-peer negotiation/delegation
Discovery	Pre-configured endpoints; no native discovery	Dynamic discovery via published Agent Cards
Typical FS application	Grounding an agent in ledgers, CRM, compliance repositories	Cross-institution/cross-vendor coordination, agentic commerce
Documented security posture	No mandatory authN/authZ in base spec; multiple 2025-2026 CVEs, tool poisoning, credential exposure	Trust depends on Agent Card verification and external identity frameworks; delegation authority not fully resolved
Payment extensibility	Not payment-native	Extended via AP2 and the A2A-x402 extension

Source: compiled by the authors from [1,2,5-10,16,31-37].

The comparison confirms the literature's central claim: MCP and A2A are architected for different layers of the agentic stack. MCP's hierarchical, pre-configured client-server model is well suited to grounding a single agent in a known, governed set of enterprise tools and data. A2A's peer-to-peer, dynamically discoverable model is designed for the harder problem of coordinating agents whose ownership, trust boundary, and even vendor identity is not fixed in advance.

B. Mapping protocols to financial-services use cases

MCP-favoured use cases share the property that a single agent, operated within one institution's trust boundary, needs governed access to internal or licensed external data and tools: connecting an LLM-based relationship-manager copilot to core banking, CRM, and market-data systems; model risk management crews reasoning over internal modeling documentation and validation records [22]; retrieval-augmented compliance research assistants requiring standardized, auditable access to internal policy repositories; and fraud-analytics copilots calling internal transaction-monitoring and case-management tools through a single, permissioned interface.

A2A-favored use cases share the property that two or more independently owned agents — potentially built by different vendors or belonging to different institutions — must discover each other's capabilities and negotiate a task: cross-institution workflows such as a corporate treasury agent negotiating with a bank-operated agent; multi-vendor orchestration within a single institution, where a fraud-detection agent, a KYC agent, and a case-management agent from three different vendors must hand off a case to one another; and agentic commerce and payments, where a consumer-facing agent must discover and transact with a merchant or bank agent it has never previously interacted with — the scenario directly targeted by AP2, A2A-x402, Mastercard's Agent Pay, and Visa's Trusted Agent Protocol [26-29].

Documented industry deployments (elaborated in Section 5) suggest that mature financial institutions are not choosing one protocol exclusively but layering them, consistent with the recommendation in the general comparative literature [7,9].

C. Security, regulatory, and liability considerations

The choice of protocol carries materially different risk profiles for a regulated institution. MCP's absence of mandatory authentication and authorization in its base specification has produced a documented and growing set of vulnerabilities — tool poisoning, credential exposure, supply-chain compromise through malicious or tampered servers — directly relevant to institutions bound by regulatory data-protection and change-management obligations [31-36]. A2A raises a different category of risk: because it is designed for cross-organizational delegation between potentially opaque agents, an institution adopting it must solve for verifying the authenticity and authority of a counterparty agent before allowing it to negotiate a task, a problem that Agent Card verification alone does not fully resolve [37]. In the payments extension of A2A, AP2 mandates, x402 settlement, and network-layer tokens from Mastercard and Visa are explicitly aimed at building the cryptographic authorization and audit trail needed to establish accountability, yet commentary consistently notes that liability frameworks for erroneous or fraudulent agent-initiated transactions remain unsettled [29,30].

This risk asymmetry connects directly to the broader financial AI-governance literature. Cross-disciplinary explainability surveys emphasize that regulatory interpretability obligations already create tension with high-performing but opaque models in finance [38,39], and comparative doctrinal analysis finds substantial jurisdictional variation in how supervisors currently address autonomous, self-directed AI systems [40]. Neither strand of literature yet addresses agent-to-agent protocols specifically, implying that MCP- and A2A-based architectures will likely need bespoke supervisory guidance rather than relying on governance frameworks built for single-model, single-decision AI systems.

D. A decision framework for financial-services architects

If the requirement is to ground a single agent in an institution's own governed tools, data, and systems, adopt MCP, paired with strict authentication, least-privilege access, and a vetted server registry to mitigate its documented security gaps.

If the requirement is to coordinate two or more independently owned or vendor-built agents — across institutions, across vendors within one institution, or in agentic commerce and payments — adopt A2A, supplemented by an external identity and authorization framework and, where payments are involved, a mandate-based extension such as AP2.

If the requirement spans both, which the literature and documented deployments suggest is the common case in mature financial-services agentic systems, deploy MCP within each individual agent for tool/data grounding and use A2A to orchestrate the resulting agent network, layering payment extensions where autonomous settlement is required.

In all cases, treat protocol adoption as a governance decision requiring sign-off from security, model-risk, and compliance functions, given the documented vulnerability surface of both protocols at their current level of maturity.

V.EMPERICAL CASE-STUDY AND COMPARATIVE ASSESSMENT

Because controlled experimental benchmarking of MCP and A2A within a live financial-services environment is not yet available in public literature, this section presents two complementary, literature-grounded empirical instruments: a case-study synthesis of documented deployments (Table 2), and a rubric-based comparative scoring exercise across six operational criteria (Table 3). Both are explicitly qualitative syntheses of cited, publicly documented evidence rather than results of a controlled experiment conducted by the authors and should be read accordingly.

VI. EMPERICAL CASE-STUDY AND COMPARATIVE ASSESSMENT

Because controlled experimental benchmarking of MCP and A2A within a live financial-services environment is not yet available in public literature, this section presents two complementary, literature-grounded empirical instruments: a case-study synthesis of documented deployments (Table 2), and a rubric-based comparative scoring exercise across six operational criteria (Table 3). Both are explicitly qualitative syntheses of cited, publicly documented evidence rather than results of a controlled experiment conducted by the authors and should be read accordingly.

A. Case-Study Synthesis of Documented Financial-Services Deployments

Table 2: Case-study synthesis of documented financial-services agentic AI deployments.

Institution / initiative	Protocol layer emphasis (as documented)	Reported outcome
JPMorgan Chase [3], [30]	MCP-style tool/data grounding across an enterprise ML platform processing ~US$10 trillion in daily transaction volume; 400+ production AI use cases	20% increase in private banking gross sales; estimated capacity to expand banker coverage by up to 50%
Goldman Sachs[3]	Agent-to-tool grounding across ~12,000 engineers paired with coding/operational agents	Reported productivity gains in accounting, compliance, and operational-finance workflows
Commonwealth Bank of Australia [3]	Homegrown agent-based fraud-detection system (s[3]ingle-institution tool grounding)	20% reduction in fraud losses, H1 FY2026
FIS with BMO / Amalgamated Bank[3]	Anthropic-powered Financial Crimes Agent (MCP-grounded compliance agent), pilot stage	Early-stage pilot deployment; outcome data not yet public
Visa [29]	Trusted Agent Protocol (A2A-adjacent, network-layer identity and consent)	Deployed to support payment-dispute resolution and corporate bill-pay workflows
Mastercard [29], [30]	Agent Pay / Agentic Tokens (A2A-adjacent, network-layer credentialing)	Global agentic-commerce network expanded to add Hong Kong
Google AP2 / A2A-x402 (Coinbase, Ethereum Foundation, MetaMask) [24], [25]	A2A payment extension for cross-vendor agent settlement	Described as production-ready for stablecoin-denominated agent-to-agent payments

Source: compiled by the authors from named institutional reporting cited in the first column

Read across rows, Table 2 shows a consistent pattern: the deployments with the longest track record and clearest quantified outcomes (JPMorgan Chase, Goldman Sachs, Commonwealth Bank) are predominantly single-institution, MCP-style tool-grounding use cases, while the deployments explicitly built around cross-organizational agent coordination (Visa, Mastercard, AP2/x402) are newer, payments-specific, and reported in terms of network expansion or production-readiness rather than quantified efficiency gains. This is consistent with A2A-style coordination being at an earlier stage of financial-services maturity than MCP-style tool grounding, while also being the layer where the highest-profile new product activity (agentic commerce) is concentrated.

B. Literature-Informed Comparative Rubric

Table 3: Literature-informed comparative scoring of MCP and A2A across six operational criteria (1 = weak, 5 = strong)

Criterion	MCP (1-5)	A2A (1-5)	Basis for score
Ease of single-agent tool/data integration	5	2	MCP was purpose-built to replace bespoke per-tool integration; A2A assumes tool access is already solved via MCP or an equivalent layer
Cross-vendor / cross-institution interoperability	2	5	MCP has no native mechanism for one agent to discover another; A2A's Agent Cards are designed exactly for this
Built-in authentication/authorization maturity (as documented, 2026)	2	3	Dozens of MCP CVEs and reports of servers with no authentication are documented; A2A depends on external identity frameworks that are still maturing
Documented production maturity in financial services	4	2	MCP-style tool grounding underlies large-scale deployments at JPMorgan and Goldman Sachs; A2A-style cross-agent coordination in FS is concentrated in payments pilots
Payment/settlement extensibility	1	5	MCP is not payment-native; A2A has direct extensions (AP2, x402) and network-layer analogues (Visa, Mastercard)
Standardization/governance neutrality	4	3	MCP governance has moved to the vendor-neutral Agentic AI Foundation; A2A remains primarily Google-led with a partner consortium

Scores reflect qualitative synthesis of the cited literature (see final column), not a controlled quantitative benchmark; see Section III.D and Section III.E for methodology and limitations.

The rubric in Table 3 quantifies the asymmetry discussed narratively in Sections IV.A-IV.C: MCP scores highest on ease of single-agent integration and current financial-services production maturity, while A2A scores highest on cross-vendor interoperability and payment extensibility. Neither protocol dominates across all six criteria, which is itself the paper's central empirical finding: a financial institution optimizing only for one criterion (for example, ease of integration) would be led toward MCP, while an institution optimizing only for cross-institutional reach (for example, agentic commerce) would be led toward A2A — but an institution facing both requirements simultaneously,

which Table 2 suggests is increasingly the norm, is led toward a layered deployment of both protocols.

VII. CONCLUSION

This paper set out to answer whether the Model Context Protocol or the Agent2Agent protocol is the right communication standard for agentic AI in financial services. The comparative architectural analysis, the literature-informed case-study synthesis, and the rubric-based comparative assessment converge on the same answer: the two protocols are not substitutes. MCP is the appropriate standard for connecting an individual financial-services agent to its tools, data, and enterprise systems, while A2A is the appropriate standard for coordinating autonomous agents across vendor, team, or institutional boundaries, including the emerging domain of agentic commerce and payments. Financial institutions such as JPMorgan Chase, Goldman Sachs, and card networks including Visa and Mastercard are already converging on layered architectures that use both protocols for different parts of the same agentic system. For financial-services leaders, the more urgent and consequential question is not which single protocol to standardize on, but how to govern the security, authentication, and liability implications of deploying both, given that neither protocol, in its current state of maturity, provides the level of built-in trust and accountability that regulated financial institutions require.

VIII. FUTURE WORK

Controlled empirical benchmarking of MCP and A2A implementations for latency, throughput, and reliability under realistic financial-services transaction volumes, extending the qualitative rubric in Table 3 into a measured benchmark.

Independent, adversarial security evaluation of MCP servers and A2A Agent Card implementations specifically within regulated financial-services deployments, extending beyond the general-purpose CVE disclosures reviewed in this paper.

Regulatory sandbox studies, conducted with central banks or financial regulators, examining liability allocation for autonomous agent-initiated payments under AP2, x402, and card-network agentic-payment frameworks.

Longitudinal tracking of standardization convergence as MCP governance moves fully into the Linux Foundation's Agentic AI Foundation and as A2A, AP2, and card-network protocols mature.

Primary case-study research, conducted directly with named financial institutions, to empirically validate the decision framework proposed in Section IV.D and the rubric in Table 3.

Integration of the explainability and algorithmic-governance literature [38-40] with agent-interoperability protocol design, to develop supervisory guidance specific to agent-to-agent financial workflows.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

REFERENCES

[1] Anthropic, "Model Context Protocol specification," 2024. Available from: https://modelcontextprotocol.io

[2] Google Research, "A2A: Agent-to-Agent protocol," 2025. Available from: https://github.com/google/A2A

[3] The CGAI Group, "Banking AI 2026: From pilots to production at scale," May 11, 2026. Available from: https://blog.thecgaigroup.com/banking-ai-2026

[4] PYMNTS, "Banks are racing into AI faster than security can follow," 2026. Available from: https://www.pymnts.com/news/artificial-intelligence/2026/banks-are-racing-into-ai-faster-than-security-can-follow/

[5] N. Yeggoli, "Model Context Protocol (MCP) vs Agent2Agent Protocol (A2A)," Medium, Apr. 21, 2025. Available from: https://medium.com/@narasimhulu.yeggoli/model-context-protocol-mcp-vs-agent2agent-protocol-a2a-896125e818c3

[6] Atlan, "MCP vs A2A protocol: Architecture, differences and when to use," May 27, 2026. Available from: https://atlan.com/know/mcp/mcp-vs-a2a-protocol/

[7] TrueFoundry, "MCP vs A2A: Compare single-agent & multi-agent protocols," Jun. 12, 2026. Available from: https://www.truefoundry.com/blog/mcp-vs-a2a

[8] Redis, "Model Context Protocol (MCP) vs. Agent2Agent (A2A): Which protocol do you need?" 2026. Available from: https://redis.io/blog/mcp-vs-a2a-which-protocol-do-you-need/

[9] Auth0, "MCP vs A2A: A guide to AI agent communication protocols," Jul. 10, 2025. Available from: https://auth0.com/blog/mcp-vs-a2a/

[10] Gravitee, "Google's Agent-to-Agent (A2A) and Anthropic's Model Context Protocol (MCP)," Jan. 20, 2026.

[11] P. D. O'Brien and R. C. Nicol, "FIPA — towards a standard for software agents," BT Technology Journal, vol. 16, no. 3, pp. 51–59, 1998. Available from: https://link.springer.com/article/10.1023/A:1009621729979

[12] Foundation for Intelligent Physical Agents, "FIPA Communicative Act Library Specification," 2002. Available from: https://www.fipa.org/specs/fipa00037/XC00037H.html

[13] J. Pitt and A. Mamdani, "Some remarks on the semantics of FIPA's agent communication language," Autonomous Agents and Multi-Agent Systems, vol. 2, no. 4, pp. 333–356, 1999. Available from: https://doi.org/10.1023/A:1010016503852

[14] M. Wooldridge and N. R. Jennings, "Intelligent agents: Theory and practice," The Knowledge Engineering Review, vol. 10, no. 2, pp. 115–152, 1995. Available from: https://tinyurl.com/2n2zsmbp

[15] Practical DevSecOps, "MCP security statistics 2026: CVEs, vulnerabilities & breach data," Jun. 26, 2026. Available from: https://dl.acm.org/doi/abs/10.1145/3814959

[16] Practical DevSecOps, "MCP security vulnerabilities: How to prevent prompt injection and tool poisoning attacks in 2026," Jan. 4, 2026. Available from: https://dl.acm.org/doi/abs/10.1145/3814959

[17] Ehtesham, A. Singh, G. K. Gupta, and S. Kumar, "A survey of agent interoperability protocols: Model Context Protocol (MCP), Agent Communication Protocol (ACP), Agent-to-Agent Protocol (A2A), and Agent Network Protocol (ANP)," arXiv:2505.02279, 2025. Available from: https://arxiv.org/abs/2505.02279

[18] P. P. Ray, "A survey on Model Context Protocol: Architecture, state-of-the-art, challenges and future directions," Authorea Preprints, 2025. Available from: https://doi.org/10.36227/techrxiv.174495492.22752319/v1

[19] K. Chang et al., "A survey of AI agent protocols," arXiv:2504.16736, 2025. Available from: https://arxiv.org/abs/2504.16736

[20] Anonymous, "Agentic artificial intelligence in finance: A comprehensive survey," arXiv:2604.21672, 2026. Available from: https://arxiv.org/abs/2604.21672

[21] Anonymous, "AI agents in financial markets: Architecture, applications, and systemic implications," arXiv:2603.13942, 2026. Available from: https://doi.org/10.3390/fintech5020034
[22] Okpala et al., "Agentic AI systems applied to tasks in financial services: Modeling and model risk management crews," 2025. Available from: https://arxiv.org/abs/2502.05439
[23] Shavandi and M. Khedmati, "A multi-agent deep reinforcement learning framework for algorithmic trading in financial markets," Expert Systems with Applications, vol. 208, p. 118124, 2022. Available from: https://doi.org/10.1016/j.eswa.2022.118124
[24] Y. Huang, C. Zhou, K. Cui, and X. Lu, "A multi-agent reinforcement learning framework for optimizing financial trading strategies based on TimesNet," Expert Systems with Applications, vol. 237, p. 121502, 2024. Available from: https://doi.org/10.1016/j.eswa.2023.121502
[25] Y. Fang, Z. Tang, K. Ren, W. Liu, L. Zhao, J. Bian, J. Li, D. Zhang, W. Yu, and Y. Liu, "Learning multi-agent intention-aware communication for optimal multi-order execution in finance," in Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2023, pp. 4003–4012. Available from: https://dl.acm.org/doi/abs/10.1145/3580305.3599856
[26] Google Cloud, "Announcing Agent Payments Protocol (AP2)," Google Cloud Blog, Sep. 16, 2025. Available from: https://ieeexplore.ieee.org/abstract/document/11395759/
[27] Coinbase, "Google Agentic Payments Protocol + x402: Agents can now actually pay each other," 2025/2026.
[28] Eco.com Support, "What is Mastercard Agent Pay? AI agent commerce protocol in 2026," Jun. 11, 2026. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7242862
[29] American Banker, "Visa, Mastercard expand agentic AI deployments," Apr. 2, 2026.
[30] CNBC, "JPMorgan Chase AI: Bank to deploy more powerful agents this year," Jun. 9, 2026. Available from: https://link.springer.com/book/10.1007/978-3-031-81647-5
[31] Aembit, "MCP security vulnerabilities: Complete guide for 2026," Mar. 19, 2026. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6972118
[32] Datastealth, "MCP security: 6 risks enterprise teams face in 2026," Apr. 30, 2026. Available from: https://datastealth.io/blogs/mcp-security
[33] UV Cyber, "Threat advisory: MCP threats," May 27, 2026. Available from: https://www.uvcyber.com/resources/reports/threat-advisory-mcp-threats
[34] Obot.ai, "MCP security risks in 2026: A CTO action plan," May 22, 2026. Available from: https://obot.ai/blog/mcp-security-cto-action-plan/
[35] S. Besen, "An unbiased comparison of MCP, ACP, and A2A protocols," Medium, Jun. 6, 2025. Available from: https://medium.com/@sandibesen/an-unbiased-comparison-of-mcp-acp-and-a2a-protocols-0b45923a20f3
[36] Crossmint, "Agentic payments protocols compared: Which is best for your AI agents? (MPP, ACP, AP2, x402)," Mar. 18, 2026. Available from: https://www.crossmint.com/learn/agentic-payments-protocols-compared
[37] Anonymous, "Permission manifests for web agents," arXiv:2601.02371, 2026. Available from: https://arxiv.org/abs/2601.02371
[38] N. Bussmann, P. Giudici, D. Marinelli, and J. Papenbrock, "Explainability for artificial intelligence in finance: A cross-disciplinary survey," Finance Research Letters, vol. 44, p. 102617, 2021. Available from: https://doi.org/10.1016/j.frl.2021.102617
[39] Černevičienė and A. Kabašinskas, "Explainable AI in finance: A systematic review," Artificial Intelligence Review, vol. 57, no. 3, pp. 1–40, 2024. Available from: https://doi.org/10.1007/s10462-024-10854-8
[40] Anonymous, "Algorithmic governance in banking: A comparative analysis of risk-based and accountability-oriented oversight," Journal of Banking Regulation, 2026. Available from: https://doi.org/10.1057/s41261-026-00320-6
[41] M. Selvam, "Adapting MCP and A2A in the Banking Ecosystem," in 2026 IEEE 5th International Conference on Computing and Machine Intelligence (ICMI), IEEE, 2026. Available from: https://doi.org/10.1109/ICMI68585.2026.11539907

ABOUT THE AUTHORS

Aditya Ghosh is a class XII student at Narayana eTechno School, Kalyan, Maharashtra, India. His research interests include Computer Science and Engineering, particularly Artificial Intelligence and Machine Learning (AI/ML), Generative AI (GenAI), and Agentic AI. His work focuses on exploring recent developments in these areas and bridging the gap between academic research and industry applications. All analyses, viewpoints, and inferences presented in this paper are expressed in his personal capacity.

Siddhartha Kumar Ghosh is Cluster Leader, Financial Services Sector, IBM Consulting, India/South Asia. The author has consulted many leading enterprises in BFSI, InsurTech, FinTech and eCommerce companies in India and abroad in AI/ML, Digital and Data driven transformation. All analysis / viewpoints expressed / inferences provided are completely on personal capacity.