



From XML Brokered Markets to Autonomous Agentic Commerce: A Design Science Framework for AI-Driven B2B Marketplaces

Michael Douglas

Millersville University, Millersville, Pennsylvania, United States

Abstract: More than two decades ago, electronic-commerce researchers proposed intelligent agent-based business-to-business marketplaces in which software agents represented buyers and suppliers within brokered exchanges. This paper revisits that brokered-marketplace concept through a design science lens and proposes an Agentic Commerce Framework in which autonomous AI agents operate within governed marketplace ecosystems. The framework integrates economic foundations, governance, trust and explainability, marketplace orchestration, specialized agents, transaction processes, controlled interoperability, and risk-adjusted human oversight. Six design principles and fifteen research propositions guide future implementation and empirical study. The paper contributes by linking electronic-marketplace theory to contemporary agentic AI, recasting the intermediary as a governance-enabled orchestrator, and establishing a research agenda for autonomous B2B commerce.

Keywords: agentic AI, autonomous commerce, B2B marketplaces, multi-agent systems, electronic markets, digital platforms, design science research

INTRODUCTION

Background

Electronic commerce has evolved through efforts to reduce search and coordination costs, improve transparency, and facilitate organizational exchange. Foundational work on intelligent agents described software entities that learn user preferences and perform delegated tasks to reduce work and information overload [1]. Electronic-marketplace research subsequently explained how digital intermediaries can reduce search costs, improve matching, and support market coordination [2]. The original brokered B2B framework placed buyer and supplier agents within a shared marketplace and used XML as the proposed messaging standard [3]. Contemporary generative and agentic AI changes the technological context through planning, memory, tool use, adaptation, and multistep execution [4, 5].

Research Problem

Existing marketplace architectures were largely developed for human-centered decisions or bounded, rule-based agents. They offer limited guidance for probabilistic, tool-using agents that initiate actions and coordinate complex commercial workflows. The challenge is to coordinate autonomous actors while preserving authority, marketplace integrity, transparency, and human control.

Research Question

How should future B2B marketplaces be designed when autonomous AI agents become active economic actors?

Contributions

This research connects foundational electronic-marketplace theory with contemporary agentic AI, introduces the Agentic Commerce Framework as a design science artifact, derives six design principles and fifteen research propositions, and establishes an empirical research agenda. Its central theoretical contribution is to conceptualize agentic commerce as a governed socio-technical platform ecosystem in which delegated machine agency, marketplace orchestration, explainability, and risk-adjusted human oversight jointly shape autonomous B2B exchange.

Theoretical Contribution

This study contributes to the information systems, electronic commerce, and agentic AI literatures by integrating previously separate research streams into a unified conceptualization of autonomous B2B commerce. Electronic marketplace research explains how digital intermediaries reduce search costs, facilitate matching, and coordinate exchange [2]. Platform ecosystem theory explains how architecture and governance shape value creation and participation [6]. Agentic AI research explains autonomous planning and execution, while delegation theory explains how authority can be distributed between people and agentic information systems [7]. These streams, however, provide limited guidance for explaining how autonomous agents, marketplace orchestration, governance, explainability, and human oversight interact within one commercial ecosystem.

The primary theoretical contribution is the conceptualization of agentic commerce as a distinct socio-technical phenomenon in which autonomous AI agents become active participants in marketplace processes rather than serving only as analytical or decision-support tools. In the Agentic Commerce Framework, specialized agents can initiate actions, coordinate workflows, support negotiation, and execute bounded commercial processes through the Agentic Marketplace Orchestrator. This formulation extends electronic marketplace theory and delegation theory into an environment characterized by distributed machine agency and cross-organizational exchange.

A second contribution is the reconceptualization of the marketplace intermediary as a governance-enabled orchestrator. The Agentic Marketplace Orchestrator does more than reduce search costs or facilitate exchange. It coordinates specialized AI agents, maintains shared transaction context, routes governance checks, manages exceptions, supports explainability, and preserves accountability. This view extends marketplace intermediation from transaction facilitation to ecosystem orchestration and embedded governance.

A third contribution is the treatment of the Governance Layer, Trust and Explainability Layer, Data and Integration Layer, and Human Oversight and Control as architectural capabilities rather than external safeguards. These cross-cutting capabilities shape agent behavior, authority boundaries, participation decisions, and organizational adoption. The framework therefore connects platform governance, explainable AI, trust,

and hybrid intelligence within a single architecture for risk-adjusted autonomous commerce [8-13].

Finally, the study contributes prescriptive and explanatory design knowledge through six design principles and fifteen research propositions. These outputs establish a structured research agenda for examining governance effectiveness, trust formation, explainability, delegated authority, marketplace mediation, multi-agent coordination, transaction outcomes, and human oversight. The Agentic Commerce Framework consequently functions as a design science artifact, a conceptual model, and a theory-building structure for future empirical investigation.

LITERATURE REVIEW

Electronic Marketplaces

Electronic marketplaces match buyers and sellers, facilitate transactions, and provide institutional infrastructure. They can reduce search costs, improve information transparency, and support more efficient matching and coordination [2]. Negotiation research further shows that electronic markets must support deal-making processes rather than merely information publication [14]. The original framework synthesized these ideas into a mediated B2B architecture [3].

Intelligent Agents in Commerce

Early intelligent-agent research envisioned software entities acting on behalf of users to reduce information overload, learn preferences, and automate recurring tasks [1]. Later commerce research applied agents to information gathering, transaction execution, monitoring, and coordination [15]. These systems established the conceptual basis for delegated software action, although their structured communication and rule-based logic limited flexible commercial reasoning.

Agentic Artificial Intelligence

Agentic AI integrates generative models with planning, persistent context, tool invocation, feedback, and multistep execution. Feuerriegel et al. [5] distinguish generative models from the broader systems and applications in which they are embedded. Bandi et al. [4] emphasize autonomy, adaptability, memory, planning, reflection, and goal pursuit. Baird and Maruping [7] establish delegation as a foundation for understanding relationships between people and agentic IS artifacts.

Multi-Agent Systems

Multi-agent systems distribute responsibilities among specialized agents that coordinate in pursuit of individual or collective goals. B2B commerce is naturally multifunctional, involving procurement, negotiation, finance, logistics, compliance, fulfillment, and support. Contemporary agents increase flexibility but strengthen the need for governance, traceability, verification, and intervention.

Human-AI Collaboration, Trust, and Explainability

Hybrid intelligence combines complementary human and machine capabilities so the combined system can perform better than either alone [9]. Trust in AI is shaped by transparency, reliability, tangibility, immediacy, and representation, although the evidence base remains sensitive to context and research design [10]. Explainable AI provides visibility into model decisions and actions and can support trustworthy use, while also raising accuracy, fairness, privacy, and stakeholder-specific explanation requirements [11, 12].

THEORETICAL FOUNDATIONS

Microeconomic Market Theory

The framework draws from market efficiency, transaction cost economics, information asymmetry, and platform economics. Electronic marketplaces create value by reducing search and coordination costs, improving transparency, and supporting matching among participants [2]. These benefits remain theoretical expectations requiring empirical testing in agentic settings.

Platform Ecosystem Theory

Agentic marketplaces coordinate interdependent participants through shared infrastructure, interfaces, and governance rules. Platform ecosystem theory emphasizes alignment among architecture, governance, and strategy [6]. Human-AI platform research further shows how human and AI capabilities can be dynamically recombined across layered modules and incorporated into platform governance [8].

Delegation and Agentic Information Systems

Delegation includes authority, accountability, monitoring, and the right to intervene. Agentic IS artifacts challenge the assumption that information systems are passive tools and create delegation relationships involving appraisal, distribution, and coordination [7]. Delegated authority may expand or contract according to trust, governance effectiveness, risk, reversibility, and oversight capacity.

Trust, Explainability, and Governance

Organizations may require observable authority limits, identifiable principals, transaction records, explanations, and override mechanisms. AI governance must address autonomy, learning, and inscrutability [13]. Transparency and reliability are important antecedents of cognitive trust [10]. XAI can support trust at both model and instance levels, but explanation depth should be aligned with stakeholder needs and transaction consequences [11, 12].

EVOLUTION FROM BROKERED MARKETS TO AGENTIC COMMERCE

The framework preserves mediation, specialization, monitoring, and shared infrastructure while adapting them to contemporary technologies and risks. XML messaging becomes one

interoperability option among APIs, agent protocols, and tool interfaces. The broker becomes a governance-enabled orchestrator. Electronic negotiation becomes agent-assisted bargaining and workflow coordination [14]. Monitoring becomes provenance, explanation, and audit, while human approval becomes risk-adjusted oversight.

RESEARCH METHOD

Design Science Research Approach

The study adopts design science research because it creates a purposeful artifact for an organizational and technological problem. The process follows problem identification, solution objectives, design, demonstration, evaluation, and communication.

Knowledge Base and Artifact Design

The artifact is informed by brokered-marketplace theory, platform ecosystem theory, intelligent and multi-agent systems, agentic AI and delegation, hybrid intelligence, trust and explainability, AI governance, and design science research. The process analyzed historical design kernels, synthesized contemporary requirements, constructed the architecture, derived design principles, and formulated propositions.

Evaluation Strategy

The present ex ante evaluation examines theoretical grounding, internal consistency, functional completeness, utility, and extensibility. It establishes conceptual plausibility rather than causal validity. Future evaluation should progress through expert review, scenario analysis, simulation, prototypes, controlled experiments, field studies, and longitudinal validation.

AGENTIC COMMERCE FRAMEWORK

The Agentic Commerce Framework is the central artifact. It extends marketplace-mediated intelligent agents into an environment characterized by agentic AI, multi-agent coordination, governance-enabled orchestration, and risk-adjusted human oversight.

Framework Overview

The layered architecture distinguishes economic purpose, governance, trust, orchestration, specialized agency, transaction execution, and outcomes. Human oversight, data, and interoperability operate across all layers.


Figure 2. Agentic Commerce Framework*A multi-layer architecture for AI-driven B2B marketplaces*

Figure 2: Agentic Commerce Framework: A multi-layer architecture for AI-driven B2B marketplaces. Source: Author.

Agentic Commerce Framework Discussion

Figure 2 presents a multi-layer architecture explaining how autonomous AI agents, platform governance, trust structures, transaction services, data integration, and human oversight collectively enable AI-driven B2B marketplaces. The framework integrates marketplace economics, platform ecosystem theory, delegated agency, explainable AI, hybrid intelligence, and responsible AI.

The Economic Foundations establish the rationale for marketplace mediation. Market efficiency, transaction cost economics, information asymmetry, and platform economics explain why organizations may adopt digital marketplaces and specialized agents. Electronic marketplaces can reduce search costs and facilitate matching and coordination, while electronic negotiation capabilities support deal making after participants are matched [2, 14]. The Governance Layer establishes the institutional and technical boundaries within which autonomous agents operate. Agent certification, policy enforcement, compliance management, audit and monitoring, and regulatory alignment constrain decision rights and preserve accountability. The layer reflects the need to align platform architecture with governance [6] and to address AI autonomy, learning, and inscrutability [13].

The Trust and Explainability Layer provides mechanisms through which participants assess autonomous decisions. Explainable AI, reputation systems, risk assessment, traceability, and transparency support informed evaluation. Trust does not arise automatically from system performance. Transparency and reliability contribute to cognitive trust, while explanations can support trust in both an overall model and a specific prediction or action [10, 11]. Explanations should also be tailored to the intended audience and purpose [12].

The Agentic Marketplace Orchestrator coordinates discovery, matching, negotiation, workflow sequencing, exception handling, and performance monitoring. Specialized buyer, supplier, financial, logistics, and compliance agents perform bounded commercial functions through the orchestrator. This design distributes delegated responsibilities rather than granting unrestricted autonomy to a monolithic system [7]. The Transaction Layer executes procurement, contract formation, fulfillment, settlement, and service and support. The Data and Integration Layer supplies enterprise data, external sources, APIs, security and privacy controls, and interoperability standards. Human Oversight and Control provides high-risk review, audit, policy configuration, intervention, override, and strategic guidance. This combination reflects hybrid intelligence, in which humans and AI contribute complementary capabilities [9].

The expected outcomes include lower transaction costs, higher market efficiency, faster procurement cycles, greater marketplace participation, improved decision support, and improved organizational performance. These are propositions to be tested, not established effects. Solid arrows indicate primary process flow, while dashed arrows indicate oversight, integration, control, and feedback.

Design Requirements

- DR1: Support autonomous commercial decision making.
- DR2: Embed governance within marketplace operations.
- DR3: Provide explainable and auditable decision processes.
- DR4: Support multi-agent coordination.
- DR5: Preserve human oversight.
- DR6: Maintain policy and regulatory compliance.

Framework Components

The six central layers are Economic Foundations, Governance, Trust and Explainability, Agentic Marketplace Orchestration, Specialized AI Agents, and Transaction Execution. Data and Integration and Human Oversight operate as cross-cutting capabilities. Together, the components preserve mediation and specialization while adding dynamic orchestration, bounded authority, controlled interoperability, risk assessment, and continuous accountability.

DESIGN PRINCIPLES

- DP1: Autonomous commercial interactions should be mediated through a shared marketplace orchestrator when common governance and accountability are required.
- DP2: Governance should be embedded within transaction processes.
- DP3: Consequential autonomous decisions should be traceable, explainable, and auditable.
- DP4: Agent authority should be bounded in proportion to risk and reversibility.
- DP5: Trust-building mechanisms should operate throughout the transaction lifecycle.
- DP6: Human intervention should remain available for high-consequence or exceptional transactions.

PROPOSITION DEVELOPMENT

Governance and Agentic Commerce

The following propositions express theoretically plausible relationships that require empirical testing.

- P1: Marketplace governance mechanisms are positively associated with organizational trust in agentic commerce systems.
- P2: The positive relationship between agent autonomy and marketplace participation is mediated by governance effectiveness.

Explainability and Autonomous Decisions

The following propositions express theoretically plausible relationships that require empirical testing.

- P3: The explainability of agent decisions is positively associated with willingness to delegate commercial authority.
- P4: Explainability positively moderates the relationship between agent autonomy and trust.

Multi-Agent Coordination

The following propositions express theoretically plausible relationships that require empirical testing.

- P5: Multi-agent architectures outperform monolithic architectures in complex B2B marketplace environments.
- P6: Agent specialization positively influences transaction effectiveness.
- P7: Agent coordination quality is positively associated with transaction efficiency.

Marketplace Mediation

The following propositions express theoretically plausible relationships that require empirical testing.

- P8: Marketplace-mediated communication reduces transaction risk relative to unrestricted peer-to-peer interaction.
- P9: Marketplace mediation positively influences perceived transaction security.

Delegated Authority

The following propositions express theoretically plausible relationships that require empirical testing.

- P10: Organizational trust positively influences delegated commercial authority.
- P11: Delegated authority positively influences transaction velocity.
- P12: Governance effectiveness positively moderates the relationship between delegated authority and organizational performance.

Human Oversight

The following propositions express theoretically plausible relationships that require empirical testing.

- P13: Human oversight positively influences trust in high-value autonomous transactions.
- P14: Autonomous-commerce effectiveness declines when transaction risk exceeds available oversight.
- P15: Human oversight positively moderates the relationship between autonomy and organizational adoption.

DISCUSSION

The framework extends marketplace theory by treating intermediation as a governance response, bridges early intelligent agents and contemporary agentic AI without claiming

equivalence, and introduces organizational delegation and accountability into multi-agent commerce. It does not propose wholesale human replacement. Instead, it specifies how human judgment, agent autonomy, marketplace controls, and explanation mechanisms can be combined within a single architecture.

Managerial and Policy Implications

Marketplace operators may need agent identity, certification, authorization, transaction records, explanation, policy enforcement, and dispute support. Procurement organizations should define authority limits, data access, escalation, and override. Autonomous commerce also complicates accountability because outcomes may emerge from multiple agents, controls, systems, data sources, and people.

FUTURE RESEARCH AGENDA

Future studies should examine autonomous procurement, inter-agent negotiation, governance and compliance, trust calibration, model-level and action-level explanations, human-AI oversight, regulation and liability, and market structure. Research designs should move beyond short-term scenarios toward prototypes, field studies, and longitudinal evidence in higher-stakes organizational contexts.

CONCLUSION

Contemporary agentic AI makes portions of the early brokered-marketplace vision more plausible while introducing governance and accountability challenges. The Agentic Commerce Framework preserves mediation, specialization, and transaction visibility while adding governance, trust, explainability, orchestration, bounded authority, controlled interoperability, and risk-adjusted oversight. It is a theoretically grounded design proposal rather than a validated operational solution.

AI USE DISCLOSURE

This manuscript was developed with assistance from Microsoft Copilot for literature exploration, conceptual organization, candidate design principles and propositions, language refinement, reference formatting, and draft visuals. AI outputs were treated as preliminary suggestions subject to human review. The author reviewed and approved the arguments, framework, propositions, interpretations, and conclusions and retains responsibility for accuracy, originality, integrity, and scholarly contribution.

REFERENCES

- [1] Maes, P. (1994). Agents that reduce work and information overload. *Communications of the ACM*, 37(7), 30-40. <https://doi.org/10.1145/176789.176792>
- [2] Bakos, Y. (1998). The emerging role of electronic marketplaces on the Internet. *Communications of the ACM*, 41(8), 35-42. <https://doi.org/10.1145/280324.280330>

- [3] Wingreen, S. C., & Douglas, M. (2000). A theory-based, brokered B2B marketplace framework with XML. Unpublished conference manuscript.
- [4] Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet*, 17(9), 404. <https://doi.org/10.3390/fi17090404>
- [5] Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111-126. <https://doi.org/10.1007/s12599-023-00834-7>
- [6] Tiwana, A. (2014). Platform ecosystems: Aligning architecture, governance, and strategy. Morgan Kaufmann. <https://doi.org/10.1016/C2012-0-06625-2>
- [7] Baird, A., & Maruping, L. M. (2021). The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. *MIS Quarterly*, 45(1), 315-341. <https://doi.org/10.25300/MISQ/2021/15882>
- [8] Rai, A., Constantinides, P., & Sarker, S. (2019). Editor's comments: Next-generation digital platforms: Toward human-AI hybrids. *MIS Quarterly*, 43(1), iii-ix. <https://doi.org/10.25300/MISQ/2019/431E0>
- [9] Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637-643. <https://doi.org/10.1007/s12599-019-00595-2>
- [10] Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627-660. <https://doi.org/10.5465/annals.2018.0057>
- [11] Rai, A. (2020). Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science*, 48, 137-141. <https://doi.org/10.1007/s11747-019-00710-5>
- [12] Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [13] Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433-1450. <https://doi.org/10.25300/MISQ/2021/16274>
- [14] Beam, C., Segev, A., Bichler, M., & Krishnan, R. (1999). On negotiations and deal making in electronic markets. *Information Systems Frontiers*, 1(3), 241-258. <https://doi.org/10.1023/A:1010050311741>
- [15] Liang, T.-P., & Huang, J.-S. (2000). A framework for applying intelligent agents to support electronic trading. *Decision Support Systems*, 28(4), 305-317. [https://doi.org/10.1016/S0167-9236\(99\)00098-6](https://doi.org/10.1016/S0167-9236(99)00098-6)
- [16] Hevner, A., vom Brocke, J., & Maedche, A. (2019). Roles of digital innovation in design science research. *Business & Information Systems Engineering*, 61(1), 3-8. <https://doi.org/10.1007/s12599-018-0571-z>