

Marketing to Machines: Delegation of Consumer Decisions to Algorithmic Agents

Milos Bujisic, and Vanja Bogicevic, NYU

Abstract— The increasing delegation of consumer purchasing decisions to autonomous algorithmic agents represents a fundamental shift in decision architectures across consumer and enterprise markets. This paper examines Marketing to Machines (M2M) as a socio-technical decision system in which Algorithmic Decision-Making Autonomous Machines (ADAM) evaluate, select, and execute purchases on behalf of humans. Using a sequential exploratory mixed-method design combining eight executive interviews and a nationally representative U.S. consumer survey ($n = 1,919$), we identify category-level variation in delegation readiness, key decision antecedents, and governance expectations shaping adoption. Results demonstrate that routinization, standardization, and low perceived risk significantly increase willingness to delegate decisions to machines, while concerns around control, transparency, and algorithmic bias act as structural constraints. We formalize M2M as a machine-centric optimization problem where marketing inputs become machine-readable signals rather than persuasive narratives. The study contributes to decision science by introducing delegation thresholds as a function of product attributes and governance trust, and offers implications for AI-driven decision systems, ethical optimization, and cyber-physical consumer infrastructures.

I. INTRODUCTION

AI-driven systems increasingly mediate economic decisions traditionally executed by humans. From auto-replenishment and subscription engines to autonomous procurement and recommendation agents, decision authority is shifting from human cognition to machine optimization. This transition challenges classical assumptions in marketing, decision theory, and cybernetics that presume a human decision-maker at the center of exchange systems.

We define Marketing to Machines (M2M) as the strategic design of products, data structures, and signals intended for evaluation by autonomous purchasing agents rather than human consumers. In this context, algorithms become decision subjects, not merely tools.

While prior research examines AI-assisted decisions, limited work addresses full delegation, where machines independently evaluate alternatives and execute transactions. This paper addresses that gap by asking:

- Which product and service categories are most susceptible to machine delegation?
- What decision antecedents govern consumer willingness to delegate?

- How do governance expectations constrain AI-driven optimization?

II. CONCEPTUAL BACKGROUND

A. From Human-Centered Persuasion to Algorithmic Delegation

Traditional marketing and decision-making models assume that humans retain final authority over purchase decisions, even when supported by recommendation systems or analytics. However, recent advances in artificial intelligence challenge this assumption by enabling Algorithmic Decision-Making Autonomous Machines (ADAM) that independently evaluate options and execute transactions. Empirical evidence shows that increasing algorithmic decision autonomy (ADA) fundamentally alters consumer responses, producing an inverted U-shaped relationship where moderate autonomy is preferred over both low and high autonomy [1].

This shift aligns with broader research on consumer autonomy in digital environments, which shows that perceived control remains central to acceptance of AI-driven decisions [8], [10]. Marketing to Machines (M2M) captures this structural transition from persuasion-oriented influence to delegation-based decision architectures, where consumers cede decision authority in exchange for efficiency and reduced cognitive effort. In this context, marketing increasingly targets machine-readable decision criteria rather than human perception.

B. Delegation Thresholds and the Calibration of AI Autonomy

Consumers are not uniformly receptive to autonomous decision-making. Low autonomy increases cognitive burden, while high autonomy threatens perceived control and reduces acceptance [1]. This creates delegation thresholds, consistent with research showing that willingness to delegate depends on perceived autonomy and decision ownership [17].

Recent studies further demonstrate that autonomy is a key mediator of AI acceptance and satisfaction [11], [15]. Systems that preserve perceived control while enabling automation achieve higher adoption. This aligns with findings that increasing user autonomy enhances recommendation acceptance [13] and reduces resistance to

algorithmic decisions [18]. Therefore, AI decision authority must be calibrated rather than maximized.

C. Algorithms as Decision Subjects in Autonomous Market Systems

While most of the literature treats algorithms as decision aids, emerging evidence shows that AI systems increasingly function as independent decision agents. Consumer Buying Agents (CBAs) autonomously select vendors, compare alternatives, and execute transactions based on predefined parameters [2].

Unlike human decision-makers, these systems prioritize structured inputs such as price, availability, and reliability, rather than symbolic or emotional attributes [2]. This reinforces earlier findings that algorithmic systems rely on structured data and rule-based evaluation [3], shifting competition toward data quality and system compatibility.

D. Machine Legibility and Data-Centric Brand Evaluation

As machine-mediated decision-making expands, brand value becomes increasingly dependent on machine legibility, defined as the extent to which product attributes are structured and interpretable by AI systems. Recommender system research demonstrates that algorithms rank and filter products based on structured attributes rather than subjective appeal [3].

Machine learning applications in marketing further reinforce this shift, enabling predictive optimization of pricing, retention, and recommendations [4], [20]. However, increased personalization introduces trade-offs between relevance and privacy, commonly referred to as the personalization–privacy paradox [21]. In M2M contexts, optimization must balance machine efficiency with governance constraints.

E. Product Category Susceptibility to Machine Delegation

Delegation varies across product categories. Routine, standardized, and low-risk purchases are more likely to be automated, as demonstrated in AI-driven grocery and smart device ecosystems [5]. These findings are consistent with broader evidence that delegation increases when tasks are objective, repetitive, and predictable [19].

Conversely, categories involving identity, values, or high consequences resist automation. Research shows that perceived autonomy and trust decline when decisions are complex or personally meaningful [16]. This reinforces the view that delegation is structurally constrained rather than universally scalable.

F. Governance, Transparency, and Explainability as Gating Constraints

Trust in AI systems depends heavily on governance mechanisms. Transparency, explainability, and regulatory compliance are critical for adoption [6], [14]. Consumers require visibility into decision processes to maintain confidence in autonomous systems.

Recent work highlights the importance of Explainable AI (XAI) and regulatory frameworks in preserving consumer autonomy [12], [14]. Blockchain-based systems and smart contracts further support trust by enabling verifiable and transparent transactions [6]. These mechanisms function as

enabling infrastructure, not optional features.

G. Bias and Risk Amplification in Delegated Decision Systems

Algorithmic systems introduce risks related to bias, manipulation, and opacity. Bias in training data can lead to systematic and scalable distortions in decision outcomes [4], [22]. Consumers are sensitive to these risks, often exhibiting algorithm aversion when systems are perceived as opaque or unfair [9].

Recent studies also highlight concerns around manipulation and loss of agency in AI-driven environments [23]. These risks reinforce the need for governance mechanisms that ensure fairness, accountability, and user control.

H. Positioning Marketing to Machines within Decision and Optimization Research

Taken together, the literature suggests that M2M represents a shift from persuasion to optimization under delegation constraints. AI systems operate on structured inputs, rule-based logic, and probabilistic optimization rather than symbolic influence.

This aligns M2M with decision science, cybernetics, and control systems, where the primary challenge is not maximizing autonomy, but designing stable, transparent, and controllable decision architectures that maintain user trust while delivering efficiency gains.

III. METHODOLOGY

This study adopts a sequential exploratory mixed-method research design to examine Marketing to Machines (M2M), defined as the delegation of consumer purchase decisions to autonomous, AI-based purchasing agents. This design is appropriate for emerging decision paradigms where conceptual boundaries are not yet fully formalized and require qualitative exploration prior to large-scale empirical testing.

The research proceeded in two phases:

- Qualitative executive interviews to identify core mechanisms, decision constraints, and category-level suitability for machine-mediated purchasing.
- Quantitative consumer survey to operationalize these mechanisms, test their generalizability, and estimate the relative strength of M2M antecedents across product categories and consumer segments.

This structure enables theory building followed by empirical validation within a decision and optimization framework.

A. Phase 1: Qualitative Executive Interviews

Eight in-depth interviews were conducted with senior executives operating at the intersection of marketing strategy, AI-enabled platforms, retail, and data-driven decision systems. Interviews followed a semi-structured protocol focused on category suitability for autonomous purchasing, conditions under which decision authority can be delegated, and key inhibitors such as trust, transparency, and loss of control.

All interviews were conducted remotely, recorded with consent, and professionally transcribed. Data were analyzed

using grounded theory methods, including open, axial, and selective coding. The objective of this phase was to identify recurring decision mechanisms and governance concerns rather than to test hypotheses. Insights from this phase informed construct definition, category selection, and survey instrument development.

B. Phase 2: Quantitative Phase

The quantitative phase consisted of a U.S. consumer survey administered through a professional research panel. The initial sample included approximately 2,500 respondents. After data-quality screening, the final analytic sample comprised 1,919 adults aged 18 and older. Quota sampling was applied to align the sample with U.S. Census benchmarks for age, gender, education, income, and geographic region.

The survey used a quasi-experimental mixed-model design. Each participant evaluated three product or service categories, randomly drawn from a fixed set of 21 categories spanning goods, services, and digital domains. For each category, respondents assessed a hypothetical AI purchasing agent acting on their behalf.

Measures captured perceived suitability of delegation, trust in the agent, perceived risks such as privacy and loss of control, and perceived benefits including time savings and decision quality. Additional measures assessed governance expectations, willingness to delegate purchasing decisions by domain, willingness to pay for AI agents, and preferred AI agent providers. Measurement items were adapted from established automation and technology-adoption scales and supplemented with new items grounded in the qualitative findings. Reliability and validity were assessed prior to substantive analysis.

Together, the two phases provide a theoretically grounded and empirically robust assessment of when, where, and why consumers are willing to delegate purchasing decisions to autonomous systems, positioning M2M as a delegated decision-making problem rather than a persuasion process.

IV. RESULTS

A. Qualitative Interview Findings

The qualitative phase consisted of eight in-depth interviews with senior executives operating at the intersection of marketing, AI-enabled platforms, and automated decision systems. Participants included:

- Antonio Lucio, Former Chief Marketing Officer, Meta (Facebook)
- Rishad Tobaccowala, Former Chief Growth Officer and Chief Strategist, Publicis Groupe
- Marc de Swaan Arons, Co-Founder, Institute for Real Growth
- Javier Meza, President, Marketing and Europe CMO, The Coca-Cola Company
- Jim Joseph, CEO, Ketchum NEXT; Former Global CEO, Wunderman Thompson
- Kathryn Metcalfe, Chief Communications Officer, CVS Health
- Kereshan Govender, CEO, Black Box Enterprises

(Blacfox)

- Yiannis Kotziagkiazouridis, Senior Leader, Choreograph Data Collective, Wunderman Thompson (WPP)

All participants had direct exposure to AI-enabled commerce, platform governance, or large-scale marketing systems undergoing automation.

Key Findings:

1. Delegation Is Conditional, Not Ideological

Executives consistently rejected the notion that consumers are either pro or anti-AI. Instead, delegation was described as situational, emerging where automation reduces friction without threatening agency. Routine, low-risk decisions were seen as natural entry points, while high-stakes or identity-relevant decisions were viewed as resistant to full automation.

2. Control and Override Are Prerequisites for Adoption

Across interviews, loss of control was identified as the primary barrier to machine-mediated purchasing. Executives emphasized that consumers are more willing to delegate when systems include visible override mechanisms, escalation paths, and reversible outcomes. Full autonomy without recourse was viewed as commercially and reputationally risky.

3. Trust Is Anchored in Governance, Not Intelligence

Participants emphasized that trust in AI purchasing agents is driven less by perceived intelligence and more by governance architecture. Transparency, explainability, and clear accountability were repeatedly cited as gating conditions for adoption. Several executives noted that performance gains alone are insufficient to justify delegation without institutional safeguards.

4. Product Category Structure Shapes Delegation Readiness

Executives converged on the view that delegation follows category-level structure. High-frequency, standardized, and outcome-predictable categories were considered well suited for autonomous purchasing. In contrast, categories involving taste, values, health, or financial consequences were described as requiring persistent human involvement.

5. Marketing Influence Shifts from Persuasion to System Design

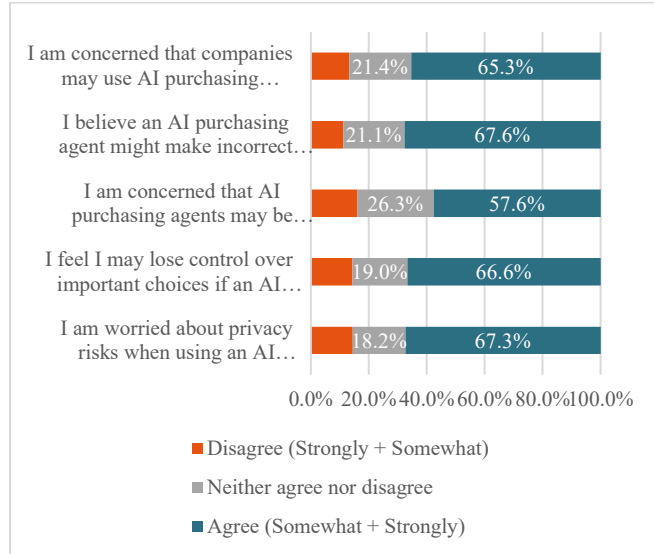
Participants consistently noted that traditional marketing levers—storytelling, emotional appeal, and brand imagery—lose effectiveness when machines execute decisions. Influence instead shifts toward data quality, attribute structure, and compatibility with algorithmic evaluation, reframing marketing as an input into decision systems rather than a persuasive act.

B. Quantitative Findings

Across the nationally representative U.S. sample (n = 1,919), consumer attitudes toward AI purchasing agents are characterized by high concern and conditional openness rather than enthusiasm. As shown in Figure 1, a majority of respondents express concern across all measured risk dimensions, including privacy, loss of control, decision accuracy, manipulation, and bias. For example, over two-


thirds of respondents report concern about privacy (67.28%), loss of control (66.65%), and incorrect decisions (67.64%). These results establish skepticism as the baseline condition for M2M adoption. Figure 1 summarizes these concerns using aggregated agreement categories (disagree / neutral / agree), providing a percentage-based view of perceived risks.

Figure 1. Customer Concern with M2M Adoption



Despite these concerns, willingness to delegate varies significantly by product category. This is shown in Figure 2, where use intention is measured on a 5-point Likert scale (1 = very unlikely to use AI purchasing agents, 5 = very likely). Observed values cluster between approximately 2.0 and 2.5 across categories. This range reflects low-to-moderate willingness, with values below 3 indicating overall reluctance to delegate decisions fully to AI systems.

Figure 2. M2M Use Intention by Product/Service Category



Higher scores (closer to 2.5) indicate relatively greater acceptance, particularly in routine and low-risk categories such as smart home devices, entertainment, and pet supplies. Lower scores (near 2.0) indicate strong resistance, particularly in high-stakes or identity-driven categories such as luxury goods, healthcare, and financial advisory services. Importantly, no category reaches the neutral midpoint (3), indicating that full delegation is not yet broadly accepted in any domain.

Neutral responses remain substantial across all categories, reinforcing that delegation is situational rather than ideological. Consumers appear willing to experiment where consequences are low or reversible, but resist delegation where decisions are perceived as consequential, value-laden, or difficult to correct.

Across categories, trust in AI purchasing agents remains low and consistently below the midpoint of the scale. Perceived suitability and comfort with delegation follow the same pattern. Even in routine categories, familiarity does not translate into confidence or comfort. These results indicate that consumers currently tolerate AI primarily as decision support, not as an autonomous decision-maker. Delegation remains fragile and contingent on safeguards rather than on performance expectations alone.

Consumers recognize clear benefits associated with AI purchasing agents, particularly time savings, effort reduction, and the ability to remember preferences over time. These efficiency-related benefits are acknowledged across most categories.


Perceived improvements in judgment or decision quality remain limited and uneven. In contrast, the prioritization of adoption factors is highly concentrated. Privacy protection and control over AI decisions clearly dominate, with 78.9% of respondents rating each as very or extremely important, far exceeding all other factors. Personalization is secondary, with 52.4% assigning high importance, while convenience ranks lowest, with 47.9% in the top category and the largest share of only moderate importance.

This distribution shows that consumers do not evaluate AI purchasing agents primarily on efficiency gains. Instead, adoption is anchored in governance guarantees, particularly the ability to retain control and protect personal data. Convenience and personalization increase attractiveness but do not compensate for weak safeguards. The pattern indicates that efficiency drives interest, but control and privacy define the threshold for actual adoption.

The results show strong resistance to paying directly for AI purchasing agents. Most respondents indicate that they would not pay for such a service, and willingness to pay clusters at very low price points. This suggests that AI purchasing agents are currently perceived as a bundled convenience feature rather than a standalone product consumers are willing to purchase explicitly. Trust in AI purchasing agents varies sharply depending on the institution controlling the system (Figure 3). Consumers show the highest willingness to delegate purchasing to AI agents developed by large,

established U.S.-based technology platforms. Other entities, including governments, non-U.S. firms, and organizations associated with surveillance or ideological influence, face substantial resistance.

Figure 3. *Trust in Entities to Implement M2M*



These findings indicate that adoption of Marketing to Machines is likely to be mediated through a narrow set of institutional gatekeepers rather than diffusing evenly across the market.

Reliability analysis showed acceptable internal consistency across constructs (Cronbach's $\alpha > 0.70$). Mean comparisons across categories were conducted, with observed differences consistent across respondents. Given the exploratory nature of the study, analysis focuses on descriptive patterns rather than causal modeling.

V. DISCUSSION AND CONCLUSION

The results position Marketing to Machines (M2M) as a delegated decision system rather than an extension of traditional consumer choice. Delegation is not driven by generalized attitudes toward AI, but emerges under specific conditions, increasing in routine, low-risk contexts and declining in high-stakes or identity-relevant domains. This confirms that delegation operates as a situational optimization strategy, where consumers trade cognitive effort for automation under bounded conditions.

These findings are consistent with prior research showing that moderate levels of algorithmic autonomy are preferred over both low and high extremes [1]. This study extends that literature by demonstrating that optimal autonomy is not fixed, but varies across product categories and decision contexts. This aligns with research showing that perceived autonomy and control influence acceptance of AI decisions

[11], [17], while adding evidence that these effects depend on the nature of the decision being delegated.

The results also reinforce prior work showing that autonomous agents prioritize structured attributes such as price, availability, and reliability [2], and that algorithmic systems rely on structured data rather than subjective evaluation [3]. However, this study shows that adoption remains constrained even when technical optimization is feasible, particularly when reversibility and control are limited.

A key contribution is the identification of governance as a gating condition. While prior research emphasizes transparency and explainability [6], [12], the findings show that privacy, control, and override capability are necessary conditions for delegation. Efficiency increases interest but does not enable adoption in the absence of trust, reinforcing concerns about bias and manipulation in AI systems [22], [23]. The findings also extend research on autonomy in AI environments by showing that autonomy functions not only as a psychological factor but as a system-level constraint. This reinforces the need for controllable AI architectures and human-in-the-loop mechanisms. At a system level, the results support a shift from persuasion-based marketing to optimization under delegation constraints, where structured data and system compatibility determine outcomes.

Future research should extend this work through behavioral experiments, cross-cultural comparisons, and formal modeling of AI decision systems as control and optimization problems.

This study has several limitations. First, the data are based on a U.S. sample, limiting generalizability across cultural contexts [8]. Second, the findings rely on hypothetical evaluations of AI purchasing agents rather than observed behavior, which may not fully reflect real-world adoption. Third, limited detail on scale development and validation constrains assessment of measurement robustness. Finally, there is a gap between executive narratives and current technological implementation, which may affect external validity.

This study shows that Marketing to Machines (M2M) is a delegated decision system in which purchasing authority is conditionally transferred to AI agents. Delegation increases in routine, low-risk contexts and declines in high-stakes or identity-relevant domains, confirming that it functions as a situational optimization strategy rather than a full transfer of agency.

The results extend prior research by showing that optimal autonomy is context-dependent and constrained by governance conditions. Privacy, control, and transparency emerge as prerequisites for adoption, while efficiency alone is insufficient. Overall, M2M is best understood as an optimization problem under delegation constraints, where adoption depends on designing controllable and trustworthy decision systems rather than maximizing automation.

REFERENCES

- [1] [1] Y. Fan and X. Liu, "Exploring the role of AI: The impact of algorithmic decision autonomy on consumer purchase decisions," *J. Bus. Res.*, vol. 145, pp. 1–13, 2022.
- [2] [2] M. Weber et al., "When algorithms go shopping: Analyzing business models for highly autonomous consumer buying agents," *MIS Q.*, vol. 44, no. 3, pp. 1–28, 2020.
- [3] [3] L. M. López-López et al., "A recommender system based on a machine learning algorithm for B2C portals," *Expert Syst. Appl.*, vol. 36, no. 2, pp. 2459–2466, 2009.
- [4] [4] V. Duarte et al., "Machine learning and marketing: A systematic literature review," *J. Bus. Res.*, vol. 141, pp. 35–52, 2022.
- [5] [5] K. Stecula et al., "Technology development in online grocery shopping," *Sustainability*, vol. 16, no. 1, 2024.
- [6] [6] N. Chotisarn and T. Phuthong, "Mapping the landscape of marketing technology," *Technol. Forecast. Soc. Change*, vol. 198, 2025.
- [7] [7] Mind Commerce, *Personal AI Market Update*, 2024.
- [8] [8] Q. André et al., "Consumer choice and autonomy in the age of AI," *Cust. Needs Solut.*, vol. 5, pp. 28–37, 2018.
- [9] [9] G. Chen et al., "Algorithm recommendation aversion," *Comput. Human Behav.*, vol. 161, 2024.
- [10] [10] R. Calvo et al., "Supporting human autonomy in AI systems," Springer, 2020.
- [11] [11] E. Cho and J. Rha, "Consumer autonomy in AI recommendations," *Int. J. Hum.-Comput. Interact.*, 2026.
- [12] [12] S. Fassiaux, "Preserving consumer autonomy through AI regulation," *Eur. J. Risk Regul.*, 2023.
- [13] [13] L. Fink et al., "Increasing user autonomy increases recommendation acceptance," *Comput. Human Behav.*, 2024.
- [14] [14] M. Mariani et al., "AI in marketing: A systematic review," *Psychol. Mark.*, vol. 39, 2022.
- [15] [15] Y. Ou and H. Li, "AI recommendation and perceived autonomy," *Curr. Psychol.*, 2025.
- [16] [16] J. Song and H. Lin, "Consumer decision delegation and trust in AI," *J. Consum. Behav.*, 2024.
- [17] [17] M. Steffel and E. Williams, "Delegating decisions," *J. Consum. Res.*, vol. 44, 2018.
- [18] [18] T. Frank and T. Otterbring, "Autonomy and AI adoption," *Br. J. Manag.*, 2024.
- [19] [19] M. Husairi and P. Rossi, "Delegation of purchasing tasks to AI," *Decis. Support Syst.*, 2024.
- [20] [20] S. Kim and J. Kim, "AI recommendation quality and satisfaction," *J. Serv. Mark.*, 2025.
- [21] [21] Y. Wang et al., "Personalization–privacy paradox," *Int. J. Inf. Manage.*, 2024.
- [22] [22] T. Zarsky, "The trouble with algorithmic decisions," *Sci. Technol. Hum. Values*, 2016.
- [23] [23] N. De Marcellis-Warin et al., "AI and consumer manipulation," *AI Ethics*, 2022.