Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Governing agentic AI in marketing: Bounded delegated agency and consumer decision authority in shopping services

Jiaying Xi

xjycastle@163.com

Zhejiang Technical Institute of Economics

Research Article

Keywords: AI shopping agents, consumer authority, delegated agency, action reversibility, goal disclosure, service governance

Posted Date: November 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11786698/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

We develop the concept of bounded delegated agency to explain how AI shopping services reduce consumer effort while preserving decision authority. Drawing on delegation, self-determination and service-value perspectives, we analyse two respondent-level datasets (N = 1,920) using factorial regression models with HC3 robust standard errors and bootstrap-based indirect association analysis. Results indicate that greater AI autonomy increases perceived task offloading, while reversibility reduces perceived agency encroachment and strengthens continued delegation intention, without demonstrated improvements in preference congruence or recommendation discrimination. Under high autonomy and low reversibility, consumer-priority disclosure is associated with an 8.91-percentage-point increase in recommendation acceptance compared with commercial-objective disclosure. Measurement summaries satisfy conventional reliability and validity thresholds. However, unavailable execution logs constrain causal interpretation and verification of actual restoration. The findings distinguish psychological reassurance from effective consumer protection, extending AI service research by demonstrating that greater acceptance and perceived control do not necessarily translate into better decision outcomes. Practically, AI service providers should establish explicit execution permissions, reliable reversal mechanisms and transparent objectives while evaluating both consumer perceptions and actual decision protection.

1. Introduction

Artificial intelligence (AI) is changing digital shopping not only by improving the quality of recommendations but also by altering who performs the actions that follow them. Conventional shopping assistants help consumers identify products, compare alternatives and evaluate offers while leaving the implementation of these decisions to the consumer. More autonomous agents can move beyond advice to modify shopping baskets, add selected items or reserve promotional offers on a customer's behalf. This transition from recommending to executing creates a consequential distinction for marketing theory. An accurate recommendation may still produce an unsatisfactory service encounter if its implementation exceeds the authority the consumer intended to delegate. Conversely, consumers may welcome considerable automation when they understand its boundaries and retain meaningful opportunities to intervene. The central issue is therefore not simply whether AI can make better shopping decisions, but under what conditions a commercial service can legitimately act for consumers without undermining their authority over the transaction. This distinction becomes increasingly important as AI-enabled commerce integrates recommendation, choice and execution within a single service interaction (Baird and Maruping, 2021; Sarstedt et al., 2026).

Research on technology-mediated services provides an important foundation for understanding this development. Service scholarship has examined how intelligent technologies redistribute activities and responsibilities among customers, frontline employees and service providers, creating new forms of interaction and value creation (Larivière et al., 2017; Wirtz et al., 2018; Huang and Rust, 2022). Complementary research on AI acceptance has identified factors associated with consumer trust,

perceived usefulness, favorable service evaluations and intentions to continue using intelligent systems (Gursoy et al., 2019; Ashfaq et al., 2020; Ameen et al., 2021). These perspectives explain why consumers may appreciate efficient, personalized and responsive services, but their explanatory reach becomes less certain when the technology is authorized to implement consequential actions. A consumer can report satisfaction with an AI assistant while remaining unaware that the system has altered a basket beyond the agreed scope. Similarly, continued-use intention may indicate confidence in an automated service without establishing whether an unwanted action could be reversed. Evaluating agentic commerce exclusively through favorable attitudes or intended usage risks conflating consumer acceptance of a service with the legitimacy and effectiveness of the actions performed on the consumer's behalf.

The difficulty becomes more apparent when delegation is examined as an allocation of decision rights rather than a general willingness to rely on technology. Delegation research emphasizes the transfer of responsibilities and authority between actors (Baird and Maruping, 2021), while studies of autonomous AI decisions examine consumer preferences for different levels of machine independence and the circumstances under which decision authority is assigned by users or initiated by technological systems (Candrian and Scherer, 2022; Adam et al., 2024). Related service research identifies perceived shared agency and distrust associated with diminished autonomy (Voyer et al., 2025; Rahman and Liang, 2026). Together, these contributions establish that consumers do not necessarily reject autonomous technologies; their responses depend partly on how control is distributed and experienced. Nevertheless, willingness to delegate, perceived autonomy and successful execution represent conceptually distinct phenomena. A consumer may voluntarily relinquish operational control over routine activities while retaining authority over financially consequential decisions. The same consumer may also regard an agent as useful while questioning whether its actions remain within an authorized boundary. What remains insufficiently explained is how specific arrangements of delegated authority shape both the psychological experience of the service encounter and the evaluation of the resulting decision.

This distinction introduces a tension between the efficiency of automated execution and the preservation of consumer authority. Greater autonomy can reduce the time, attention and repeated confirmation required to complete shopping tasks. Yet reducing consumer involvement may also make it difficult to determine when an advisory interaction has become an executable commitment. Consider an AI agent permitted to assemble a basket but not to finalize a purchase. Adding a suitable item may be acceptable, whereas securing an offer or proceeding towards payment without the required authorization may be perceived as overstepping, even when the product matches the consumer's preferences. Conversely, an agent that repeatedly seeks confirmation for inconsequential actions may preserve visible decision control while diminishing the convenience that motivated delegation. These examples suggest that neither maximum automation nor maximum user intervention necessarily constitutes an optimal service arrangement. The relevant question concerns the alignment between the authority granted to the agent and the authority retained by the consumer. This alignment is important because a service can deliver a desirable consumption outcome while violating the procedural expectations governing how that outcome should be achieved.

Reversibility adds a further dimension to this problem because authorization and recovery are not equivalent forms of consumer protection. An agent may operate within an initially accepted scope yet produce an outcome that the consumer subsequently wishes to change. Providing an accessible undo function can reduce anticipated vulnerability by signaling that delegated actions remain subject to consumer intervention. However, the availability of a reversal option does not necessarily establish that the original transaction state can be restored. Promotional offers may expire, inventory may change or reservations may become unavailable between execution and attempted correction. Consequently, perceived reversibility and effective restoration should not be treated as interchangeable indicators of consumer control. The distinction also complicates the relationship between autonomy and acceptance. Consumers may tolerate more extensive delegation when they believe unwanted actions can be corrected, but this reassurance may not correspond to equivalent operational protection. Conversely, limited reversibility may make autonomous execution more consequential even when the agent's recommendation is appropriate. Understanding this tension requires attention to both the psychological reassurance associated with retained intervention rights and the practical limitations of post-execution recovery.

A related issue concerns the objectives that guide an agent's actions. AI shopping services may be expected to prioritize consumer preferences, but the broader commercial environment can introduce competing considerations involving promotional visibility, conversion targets or retailer-defined priorities. Recent research has incorporated stakeholder ethics into agentic delegation and examined how machine initiative relates to interaction value (Saeed and Prybutok, 2026; Wang et al., 2026). These developments highlight the importance of communicating whose interests are represented when automated systems participate in commercial decisions. Nevertheless, disclosure should not be equated with either comprehension or preference alignment. Informing consumers that an agent follows a particular objective may improve the interpretability of its behavior without ensuring that the objective is understood, accepted or consistent with the consumer's priorities. Furthermore, a consumer who recognizes an unwanted objective and declines a recommendation may be exercising effective authority rather than demonstrating service failure. Recommendation acceptance, therefore, cannot independently establish whether a delegated arrangement respects consumer interests. The relationship among communicated objectives, perceived control and behavioral responses must be examined without assuming that favorable choices necessarily reflect appropriate governance.

To address these issues, we introduce bounded delegated agency, defined as an arrangement in which an AI service receives a specified scope of execution authority while consumers retain identifiable intervention rights and are informed of the priorities guiding delegated actions. The concept distinguishes three elements that are often discussed together but perform different functions: execution permission determines what an agent may do, reversibility concerns opportunities to correct its actions, and objective disclosure communicates the basis on which delegated decisions are made. Importantly, these elements need not produce identical psychological or behavioral responses. Permission may reduce interaction burden while increasing perceived encroachment; reversibility may provide reassurance without guaranteeing restoration; and disclosure may support informed judgment

without necessarily increasing recommendation acceptance. We use agency encroachment to describe the perceived crossing of a consumer-retained decision boundary, rather than dissatisfaction with the substantive quality of a recommendation. This distinction separates the appropriateness of an action from the desirability of its outcome and provides a basis for examining how consumers evaluate different configurations of delegated service authority.

Our research consequently shifts attention from the general acceptance of autonomous AI to the conditions under which execution authority is allocated, experienced and evaluated in consumer transactions. The theoretical contribution lies in treating consumer authority as a structured relationship between delegated capabilities and retained rights, rather than as the inverse of technological autonomy. By distinguishing reductions in decision burden from perceived protection against unwanted actions, the framework explains why favorable service evaluations and concerns about agency encroachment may coexist. It also clarifies why consumers' responses to recommendations should be interpreted alongside the conditions governing execution, rather than treated as sufficient evidence of a successful delegation arrangement. For marketing theory, this perspective extends service-role allocation and AI acceptance research towards an account of how commercially consequential actions become authorized and contestable. For service design, it directs attention to the boundaries of permitted execution, the credibility of corrective options and the communication of decision priorities. More broadly, it positions the preservation of consumer authority as a distinct criterion for evaluating AI-mediated service relationships, alongside convenience, recommendation quality and consumer acceptance.

2. Literature review and hypothesis development

2.1 Theoretical underpinnings

Delegation transfers task rights and responsibilities (Baird and Maruping, 2021). Shopping permission may cover comparison, basket construction or reservation while excluding payment. Execution scope differs from invocation and appearance: a user-invoked agent can have extensive permission, whereas an automatically displayed assistant may only advise (Araujo, 2018; Adam et al., 2024).

Self-determination theory explains why delegation can preserve autonomy. Volition concerns endorsement and ownership, including choosing another actor to perform an action (Ryan and Deci, 2020). This psychological endorsement does not establish that a service implements the agreed boundary. A consumer may endorse a weakly protected service or dislike a cumbersome safeguard that nevertheless prevents an unwanted commitment.

Service logic emphasizes consequences in use and interaction (Grönroos, 2008; Vargo et al., 2008; Ramaswamy and Ozcan, 2018). Reducing effort and preserving an intended choice are distinct potential benefits of delegation. We connect execution rights to both outcomes without assuming that either compensates for failure on the other.

We define bounded delegated agency through permitted execution, restoration options and disclosed priorities. This service configuration is conceptually distinct from a higher-order attitude. Table I separates these design features from their psychological and behavioral consequences. Reversibility concerns recovery from a past action; revocation withdraws permission for future execution. A service may offer one without the other. Disclosure makes a stated objective interpretable but cannot, by itself, verify that the ranking rule follows it.

Table I. Constructs and operational distinctions

Construct	Meaning	Operational indicator
Action autonomy A	Execution within a permission scope without stepwise confirmation	Condition code; separate invocation check
Reversibility R	Restoration of a previous state after action	Condition code; access and successful recovery
Goal disclosure G	Stated consumer priority or commercial objective	Disclosure code; comprehension; independent rule audit
Task offloading	Perceived reduction in processing and operational burden	OFF1–OFF4
Agency encroachment	Perceived crossing of a retained decision boundary	ENC1–ENC4; distinguish neighboring constructs
Control and trust	Distinct evaluations of intervention and reliance	CTL1–CTL3; TRU1–TRU3
Continued delegation	Intention to use and authorize similar assistance	CON1–CON3
Preference congruence	Final choice agreement with specified task preferences	Task-specific index from 0 to 1
Recommendation acceptance	Proportion of presented recommendations accepted	Accepted recommendations / trials (0–1)
Selective acceptance	Matched less mismatched recommendation acceptance	Within-respondent difference in Study 2

Note. A = action autonomy; R = reversibility; G = disclosed objective. Restoration and future revocation are distinct rights. Design codes, perceptions and behavioral indicators refer to different stages of the service process. Source: Authors' conceptualization.

Authority attaches to a particular action and stage. Permission to compare cannot establish permission to reserve or pay. A general control rating may conceal these different rights. Stakeholder interests are already central to ethical delegation (Saeed and Prybutok, 2026); our framework specifies how retained intervention rights relate to measurable shopping outcomes.

2.2 Task offloading and continued delegation

Cognitive offloading reduces internally maintained processing demands (Risko and Gilbert, 2016). An executing shopping agent can remove repeated comparison, memory and implementation work. Utilitarian assistant use and algorithm appreciation motivate this benefit route, while leaving execution permission as a separate question (McLean and Osei-Frimpong, 2019; Logg et al., 2019).

Delegation also requires supervision: consumers must understand permissions, inspect execution and repair unwanted actions. These demands may absorb the work removed. A valid comparison therefore holds decision requirements and recommendation information constant while varying implementation. Supervision costs should be specified independently through procedures or observed effort.

Experienced relief can encourage continued delegation by preserving expected convenience. This benefit can coexist with encroachment if an agent removes work while crossing a retained boundary. We propose parallel routes to explain why a favorable efficiency evaluation need not produce equally favorable continuation intentions.

H1. When supervision does not offset the work removed, higher AI action autonomy increases perceived task offloading, which is positively associated with continued delegation intention.

2.3 Agency encroachment and corrective control

Agency encroachment is the perceived crossing of a retained decision boundary. Adding a suitable product prematurely may create encroachment; presenting unsuitable advice for voluntary rejection need not. Control and trust are neighboring explanations (Pizzi et al., 2021; Frank et al., 2023). We therefore distinguish responses to execution beyond permission from generalized dissatisfaction or distrust.

Holding the executed option and its quality constant, changing understood permission should change encroachment if boundary crossing is the mechanism. If responses instead follow quality irrespective of permission, dissatisfaction is a simpler account. Uncertain boundaries and difficult correction expose consumers to unintended execution, motivating H2.

H2. Where the execution boundary remains uncertain, higher AI action autonomy under low reversibility increases perceived agency encroachment, which is negatively associated with continued delegation intention.

Reversibility can reduce the anticipated cost of unwanted execution before recovery occurs. This reassurance differs from changing the original permission: restoration cannot retrospectively authorize an action outside scope. H3 therefore predicts attenuated encroachment rather than complete removal of authority concerns.

H3. Higher reversibility attenuates the positive relationship between AI action autonomy and agency encroachment by reducing the anticipated cost of unwanted execution.

Actual protection requires detectable mismatch, an accessible intervention and a preferable feasible alternative. Undo cannot restore expired stock or a lost price; an already optimal selection leaves little

room for improvement. H4 concerns final congruence where correction is possible. These conditions require specification before assignment, rather than selecting cases based on successful correction.

H4. When preference mismatch is observable and a preferable alternative remains available, higher reversibility improves the effect of AI action autonomy on final preference congruence.

2.4 Disclosed objectives and recommendation acceptance

Provider objectives create an agency problem distinct from execution scope (Jensen and Meckling, 1976). Permission to act may coexist with misunderstanding whether ranking prioritizes consumer preferences or commercial interests. Ethical design and explanation research distinguish communication, comprehension and alignment (Du and Xie, 2021; Shin, 2021; Nizette et al., 2025). Disclosure can change beliefs without changing the selection rule.

Under high autonomy and limited correction, consumers depend on the agent's execution before they can repair an unwanted choice. A stated consumer priority can make accepting a recommendation more attractive than a commercial objective. We therefore predict a higher recommendation acceptance rate following consumer-priority disclosure in this condition. This prediction concerns the disclosed objective; establishing whether comprehension or endorsement explains the response requires separate evidence.

H5. Under high AI action autonomy and low reversibility, consumer-priority goal disclosure increases the rate of recommendation acceptance relative to commercial-objective disclosure.

We define recommendation acceptance as the proportion of presented recommendations that a consumer accepts. It differs from continued delegation intention and selective acceptance. Consumers can accept a suggestion yet subsequently reject purchase, and overall acceptance can increase for both matched and mismatched options. H5 concerns uptake; H6 separately addresses recommendation discrimination.

2.5 Recommendation discrimination and service value

Service interactions can create or destroy value (Plé and Chumpitaz Cáceres, 2010; Järvi et al., 2018; Čaić et al., 2018; Chandra and Rahman, 2024). Ceding control differs from appropriate reliance (Ivanova-Stenzel and Tolksdorf, 2024; Lee and See, 2004). Selective acceptance is matched minus mismatched acceptance. Overall acceptance can rise without improving this difference.

Reversibility may encourage exploration while preserving rejection of unsuitable options. It may instead reduce inspection because errors appear repairable. H6 predicts a net improvement in selectivity under comparable opportunities to recognize mismatch. That prediction can fail even if reversibility provides psychological reassurance.

H6. Under comparable opportunities to recognize recommendation mismatch, higher reversibility increases the difference between acceptance rates for preference-matched and preference-mismatched

recommendations.

Acceptance, restoration and retention occur at different stages. Removing an initially accepted recommendation may improve final congruence without changing initial discrimination. In Fig. 1, we distinguish these behavioral criteria from the psychological routes through offloading and encroachment.

Note

A = action autonomy; R = reversibility; G = goal disclosure. H5 compares consumer-priority and commercial-objective disclosure on recommendation acceptance at A = 1, R = 0. Source: Authors' conceptualization.

3. Methodology

3.1 Research approach and analytical units

We develop our framework through theory adaptation and conceptual modeling (Jaakkola, 2020) and assess its observable implications in two deidentified datasets containing 640 and 1,280 respondents. We analyse condition codes, demographics, composite measures and task outcomes. Table II describes the factorial structure and sample profiles. We retain all available respondent records; identifiers are unique within each study.

Study 1 specifies one recommendation decision per respondent; Study 2 specifies four, comprising two matched and two mismatched recommendations. We analyse 1,280 Study 2 respondent summaries, without treating the 5,120 implied decisions as independent observations. Event logs, task order and documentation sufficient to verify recruitment, allocation and consent are unavailable in our analytical files. We therefore interpret contrasts as associations.

Table II. Factorial structure and respondent profiles

Feature	Study 1	Study 2
Factorial structure	A × R	A × R × G
Respondent records	640	1,280
Cells and records per cell	4 × 160	8 × 160
Goal disclosure	G = 1 fixed	G = 0 or 1
Decisions per respondent in design	1	4
Individual event logs available	No	No
Age mean (SD)	36.65 (11.83)	37.06 (12.14)
Female n (%)	326 (50.9%)	663 (51.8%)
Never used generative AI n (%)	122 (19.1%)	205 (16.0%)
Recorded macro-regions	4	4

Note. Study 2 specifies two matched and two mismatched decisions per respondent. Female counts exclude other or undisclosed gender categories. Cell sizes describe the analytical sample; equal cell sizes alone do not establish random allocation. Source: Authors' analyses.

3.2 Design and measures

We analyse a balanced A × R structure in Study 1 and an A × R × G structure in Study 2, with 160 respondents per cell. We code consumer implementation of advice as A = 0 and agent execution within the specified scope as A = 1. Ordinary correction is R = 0 and accessible restoration is R = 1. Commercial-objective disclosure is G = 0 and consumer-priority disclosure is G = 1; G is fixed at 1 in Study 1. These codes identify the service arrangement. Our analysis does not separately estimate stimulus comprehension or manipulation strength.

We analyse five-point measures of task offloading (OFF, four items), agency encroachment (ENC, four items), perceived control (CTL, three items), trust (TRU, three items) and continued delegation intention (CON, three items). We calculate composite means and standard deviations and summarize Cronbach's alpha, composite reliability (CR), average variance extracted (AVE), standardized loadings and heterotrait–monotrait (HTMT) ratios from the measurement records. We interpret AVE and CR under a standardized reflective model (Fornell and Larcker, 1981), and HTMT as construct separation (Henseler et al., 2015). Table III presents reliability, AVE and minimum loadings; Table IV presents the complete HTMT matrices for both studies. Supplementary Table S1 gives individual item loadings. Without respondent-linked items or complete model output, we do not re-estimate the measurement model.

We analyse task duration, overall recommendation acceptance, preference congruence and, in Study 2, the matched–mismatched acceptance difference. We define overall acceptance as accepted recommendations divided by trials, following the data dictionary. H5 uses this observed outcome. We changed H5's outcome from purchase authorization to recommendation acceptance during analysis and

therefore evaluate H5 as exploratory. Neither this outcome change nor the seven-comparison testing family was preregistered. Eq. 1 defines the preference-congruence index.

$$PC_i = \frac{U_i\,(p\,chosen)\ -\ minU_i\,(p)}{maxU_i\,(p)\ -\ minU_i\,(p)}\ (1)$$

The task utility score combines normalized price, battery duration and warranty with a fixed 0.34 budget-exceedance penalty. In Eq. 1, we rescale chosen utility over the feasible range and assign 1 when utilities are tied. This task-specific index measures preference congruence, rather than willingness to pay or welfare. We use the recorded index because the analytical files do not contain the choices, budgets and weights required to reconstruct it.

In Eq. 2, we define respondent i's discrimination score, D, as matched minus mismatched acceptance. With two recommendations in each class and equal class weights, we recover the matched acceptance rate as overall acceptance + D/2 and the mismatched rate as overall acceptance − D/2. Both derived rates take admissible values of 0, 0.5 or 1 for all 1,280 Study 2 respondents. These identities recover class-level rates; they do not reconstruct event order or individual trial decisions. Study 1 does not measure discrimination.

$$D_i = \ Pr\,(accept \mid\ matched,\ i)\ -\ Pr(accept \mid mismatched,\ i) \quad (2)$$

3.3 Statistical specification

We estimate saturated factorial regressions with binary coding (Eq. 3). In Study 1, we omit the G terms, so the A coefficient represents the autonomy contrast at R = 0. In Study 2, R = G = 0 is the reference. We average contrasts with equal weights across other factor levels; the average AR contrast in Study 2 is βAR + 0.5βARG. We distinguish conditional and average contrasts throughout.

To test H5, we estimate βG + βAG at A = 1, R = 0 in Study 2, comparing two cells totaling 320 respondents. We assess sensitivity by adding demographic covariates and by resampling respondents within these two cells 10,000 times (seed 20261110). We retain the respondent as the analytical unit.

$$Y_i = \beta_{\,0} + \beta_{\,A} A_i + \beta_{\,R} R_i + \beta_{\,G} G_i + \beta_{\,AR} A_i R_i$$

$$+\beta_{\,AG} A_i G_i + \beta_{\,RG} R_i G_i + \beta_{\,ARG} A_i R_i G_i + \epsilon_{\,i}\ (3)$$

We estimate HC3 covariance and t intervals using residual degrees of freedom (MacKinnon and White, 1985). We analyse Study 2 discrimination as one paired difference per respondent (Cameron and Miller, 2015), with a one-sample t interval for its mean. HC3 accommodates heteroskedasticity under respondent independence. Without session identifiers or event logs, we cannot adjust for clustering across respondents or event order.

We estimate conditional indirect associations for mediator m at reversibility r and objective g (Eq. 4). Our parallel-mediator model includes all factorial terms and constant mediator–outcome slopes, omitting G terms in Study 1. We jointly resample respondents and their variables within cells 5,000 times (seeds 20261016 and 20261017), using the same draws for high-minus-low reversibility differences in indirect products. Percentile intervals describe statistical associations; causal mediation requires additional identification assumptions (Imai et al., 2010).

$$IE_m(r, g) = [a_{mA} + a_{mAR}r + a_{mAG}g + a_{mARG}rg]b_m \quad (4)$$

We apply Holm correction to seven comparisons: the six A × R contrasts for encroachment, continuation and congruence across both studies, and the H5 disclosure contrast (Holm, 1979). We report unadjusted inference for the remaining acceptance, offloading, duration, discrimination and indirect-product comparisons. The correction controls multiplicity within this stated family, rather than across every reported outcome. In covariate sensitivity analyses, we adjust for age, gender, region and AI-use frequency.

4. Results

We report measurement summaries, psychological responses and behavioral outcomes. Composite measures use five-point units; congruence and acceptance range from 0 to 1. Discrimination is a within-respondent difference in acceptance rates. We distinguish conditional from average contrasts and report Holm-adjusted p values for the seven comparisons specified in Section 3.3.

4.1 Sample and measurement summaries

We retain complete observations on the composite measures, conditions and principal task outcomes. Each factorial cell contains 160 respondents, with ages ranging from 18 to 65. Table II presents sample profiles, and Table III combines the composite means and standard deviations with the measurement indices.

Table III. Composite measures and measurement indices

Study	Scale	M	SD	α	CR	AVE	Min λ
1	OFF	3.182	0.770	0.8508	0.8510	0.5882	0.7415
1	ENC	2.854	0.866	0.8689	0.8690	0.6238	0.7710
1	CTL	3.351	0.787	0.7972	0.7976	0.5679	0.7240
1	TRU	3.487	0.781	0.7882	0.7891	0.5556	0.7194
1	CON	3.441	0.798	0.8123	0.8125	0.5909	0.7640
2	OFF	3.181	0.775	0.8359	0.8360	0.5603	0.7415
2	ENC	2.989	0.858	0.8650	0.8651	0.6158	0.7650
2	CTL	3.278	0.808	0.8016	0.8017	0.5741	0.7518
2	TRU	3.346	0.795	0.7950	0.7952	0.5643	0.7248
2	CON	3.366	0.796	0.7959	0.7966	0.5667	0.7107

Note. n = 640 in Study 1 and 1,280 in Study 2. OFF = task offloading; ENC = agency encroachment; CTL = perceived control; TRU = trust; CON = continued delegation intention. M = mean; SD = standard deviation; α = Cronbach's alpha; CR = composite reliability; AVE = average variance extracted; Min λ = minimum standardized loading. We retain four decimals from the measurement records and do not re-estimate these indices from respondent-linked items. Individual item loadings appear in supplementary Table S1; the full HTMT matrices appear in Table IV. Source: Authors' descriptive statistics and accompanying measurement records.

Composite means range from 2.854 to 3.487 in Study 1 and from 2.989 to 3.366 in Study 2, with standard deviations of 0.770–0.866 across the two studies.

The reliability pattern is consistent across the two studies. In Study 1, Cronbach's alpha ranges from 0.7882 to 0.8689 and CR from 0.7891 to 0.8690. In Study 2, the corresponding ranges are 0.7950–0.8650 and 0.7952–0.8651. Trust has the lowest alpha and CR in both studies, while encroachment has the highest. The coefficients support the internal consistency of all five composites, including the three-item control, trust and continuation measures. This consistency concerns how well the items within each scale hold together; whether neighboring scales capture distinguishable constructs requires the separate HTMT assessment below.

Convergent validity follows the same broad pattern. AVE ranges from 0.5556 to 0.6238 in Study 1 and from 0.5603 to 0.6158 in Study 2, exceeding the 0.50 reference point (Fornell and Larcker, 1981). Under the standardized reflective model, these values indicate that each construct accounts, on average, for more than half of its indicators' variance. The minimum standardized loading for each construct ranges from 0.7194 to 0.7710 in Study 1 and from 0.7107 to 0.7650 in Study 2; every minimum exceeds 0.70. The agreement between AVE and the loading minima supports convergence within each scale, alongside the reliability evidence in Table III.

Table IV presents all ten unique HTMT comparisons in each study. The ratios range from 0.035 to 0.362 in Study 1 and from 0.033 to 0.408 in Study 2. Every pair falls below the conservative 0.85 criterion (Henseler et al., 2015). The largest ratio concerns encroachment and trust in Study 1 (0.362081 before rounding), and encroachment and perceived control in Study 2 (0.408422). Thus, even the closest measured neighbors of encroachment remain below the criterion. Reporting both matrices also makes clear that this assessment covers all five constructs, rather than only the two proposed mediators.

Table IV. Discriminant validity across the two studies

Panel A Study 1 (n = 640)					
Construct	**OFF**	**ENC**	**CTL**	**TRU**	**CON**
OFF	—	—	—	—	—
ENC	0.035	—	—	—	—
CTL	0.066	0.289	—	—	—
TRU	0.223	0.362	0.129	—	—
CON	0.276	0.253	0.069	0.093	—
Panel B Study 2 (n = 1,280)					
Construct	**OFF**	**ENC**	**CTL**	**TRU**	**CON**
OFF	—	—	—	—	—
ENC	0.033	—	—	—	—
CTL	0.060	0.408	—	—	—
TRU	0.234	0.345	0.148	—	—
CON	0.280	0.242	0.108	0.129	—

Note. Panel A: Study 1, n = 640; Panel B: Study 2, n = 1,280. OFF = task offloading; ENC = agency encroachment; CTL = perceived control; TRU = trust; CON = continued delegation intention. Lower triangles show unique construct pairs; diagonal and duplicated entries are omitted. Values are rounded to three decimals. Maximum unrounded off-diagonal HTMT = 0.362081 in Study 1 and 0.408422 in Study 2, both below 0.85. These are point-estimate comparisons; bootstrap confidence intervals and an inferential test against unity are unavailable. Source: Accompanying measurement records.

The offloading–encroachment ratios are particularly low: 0.035 in Study 1 and 0.033 in Study 2. This pattern is consistent with distinguishing relief from task demands from perceived crossing of a decision boundary, as our framework proposes. Continued delegation is also distinct from both routes: its HTMT ratios with offloading are 0.276 and 0.280, and those with encroachment are 0.253 and 0.242. Favorable offloading ratings, encroachment and willingness to continue therefore remain separate measured

evaluations. Their separation permits us to examine the two routes to continuation without treating the outcome as another indicator of either mediator.

We interpret these results as evidence of internal consistency, convergent validity and discriminant validity under the HTMT point-estimate criterion. We do not treat the matrices as bootstrap tests against unity because the requisite intervals are unavailable. Nor do they establish global measurement-model fit, equivalence across conditions or separation from unmeasured reactance and dissatisfaction. The subsequent models consequently concern associations among observed composites. A low HTMT ratio does not itself demonstrate that the corresponding psychological processes are causally independent.

4.2 Task offloading and the adverse encroachment route

We estimate an autonomy contrast in offloading of 0.211 scale points [95% CI 0.043, 0.379], p = 0.014, at low reversibility in Study 1 (Table V). In Study 2, the contrast averaged equally over reversibility and objective disclosure is 0.272 [0.188, 0.356], p < 0.001 (Table VI). Both estimates support the offloading component of H1, although they concern different conditional comparisons.

We estimate average autonomy contrasts in task duration of − 22.13 seconds [− 29.78, − 14.48] and − 22.48 [− 27.82, − 17.15] in Studies 1 and 2. These secondary results complement perceived offloading. We cannot infer total delegation costs because the event logs needed to establish whether duration includes review and repair are unavailable.

In Study 1, we estimate offloading indirect associations with continued delegation of 0.047 [0.009, 0.092] at low reversibility and 0.088 [0.043, 0.144] at high reversibility (Table VII). All four conditional offloading products in Study 2 are positive. We interpret this pattern as consistent with H1’s benefit route. Our mediation estimates remain observational because the mediator–outcome relationship is not randomized and supervision costs are not separately measured.

We also observe the adverse route through encroachment. In Study 1, the autonomy contrast at low reversibility is 0.428 [0.242, 0.614], and the corresponding indirect association with continuation is − 0.080 [− 0.132, − 0.039]. In Study 2, encroachment products are negative under both objectives at low reversibility. These associations are consistent with H2. We cannot isolate the permission-boundary mechanism without varying or measuring the retained boundary independently of product quality.

Table V. Study 1 factorial contrasts

Outcome and contrast	Estimate	95% CI	p	Holm p
Offloading A at R0	0.211	[0.043, 0.379]	0.014	—
Encroachment A at R0	0.428	[0.242, 0.614]	< 0.001	—
Encroachment A at R1	0.005	[-0.181, 0.190]	0.960	—
Encroachment AR	-0.423	[-0.686, -0.161]	0.002	0.006
Continuation A at R0	-0.027	[-0.203, 0.149]	0.762	—
Continuation A at R1	0.342	[0.171, 0.512]	< 0.001	—
Continuation AR	0.369	[0.124, 0.613]	0.003	0.010
Congruence AR	-0.009	[-0.073, 0.055]	0.784	0.784
Acceptance AR	0.063	[-0.085, 0.210]	0.405	—
Duration A average	-22.130	[-29.779, -14.481]	< 0.001	—

Note. n = 640 respondents. A at R0/R1 denotes a conditional autonomy contrast; AR denotes the difference across reversibility levels. Duration is in seconds; other outcomes use scale or 0–1 units. HC3 t intervals are unadjusted. Holm p values cover the six interactions and H5 across Tables V and VI. "—" denotes a comparison outside this seven-test family. Source: Authors' analyses.

Table VI. Study 2 factorial contrasts

Outcome and contrast	Estimate	95% CI	p	Holm p
Offloading A average	0.272	[0.188, 0.356]	< 0.001	—
Encroachment AR average	-0.416	[-0.596, -0.237]	< 0.001	< 0.001
Continuation AR average	0.312	[0.141, 0.484]	< 0.001	0.002
Congruence AR average	-0.028	[-0.056, -0.001]	0.040	0.080
Acceptance AR average	0.075	[0.022, 0.128]	0.005	—
Acceptance G at A1 R0 (H5)	0.089	[0.035, 0.143]	0.001	0.006
Trust G average	0.156	[0.070, 0.243]	< 0.001	—
Encroachment ARG	0.267	[-0.091, 0.625]	0.144	—
Duration A average	-22.481	[-27.815, -17.147]	< 0.001	—

Note. n = 1,280 respondent records. Average AR = βAR + 0.5βARG. G compares consumer-priority with commercial-objective disclosure. The H5 comparison uses two cells with 160 respondents each. HC3 t intervals are unadjusted; Holm p values cover the seven-test family specified in Section 3.3. "—" denotes another comparison. Source: Authors' respondent-level analyses.

4.3 Reversibility and psychological responses

The encroachment A × R interaction is -0.423 [-0.686, -0.161] in Study 1 and − 0.416 [-0.596, -0.237] in Study 2. Holm-adjusted p values are 0.006 and < 0.001. In Study 1, the autonomy contrast falls from 0.428 at low reversibility to 0.005 at high reversibility. The interaction describes this change in the autonomy contrast, rather than the ranking of all four condition means.

Study 1's high-reversibility interval, [− 0.181, 0.190], is compatible with attenuation but does not establish equivalence or elimination. Study 2 averages G equally. Its three-way encroachment contrast is 0.267 [− 0.091, 0.625], p = 0.144, providing no established objective-related heterogeneity in attenuation.

Continued delegation has positive A × R interactions of 0.369 [0.124, 0.613] and 0.312 [0.141, 0.484], with Holm-adjusted p = 0.010 and 0.002. Together with the encroachment results, these contrasts support the psychological pattern proposed in H3. Figure 2 places these estimates beside the behavioral contrasts. The association does not identify anticipated recovery cost as the causal mechanism, and the continuation outcome remains an intention rather than observed repeat use.

We test reversibility differences in indirect products directly (Table VII). Encroachment products become less negative at high reversibility in Study 1 and under both Study 2 objectives; the unadjusted bootstrap intervals for these differences lie above zero. Intervals for the offloading-product differences include zero. We therefore find clearer evidence of moderation through attenuation of the adverse route than through strengthening of the benefit route. This comparison uses intervals for the differences themselves.

Table VII. Conditional indirect associations and reversibility differences

Study	Condition	Via offloading Estimate [95% CI]	Via encroachment Estimate [95% CI]
1	R = 0; G = 1	0.047 [0.009, 0.092]	-0.080 [-0.132, -0.039]
1	R = 1; G = 1	0.088 [0.043, 0.144]	-0.001 [-0.038, 0.033]
1	ΔR; G = 1	0.041 [-0.010, 0.102]	0.079 [0.027, 0.143]
2	R = 0; G = 0	0.043 [0.009, 0.081]	-0.108 [-0.156, -0.066]
2	R = 0; G = 1	0.068 [0.030, 0.112]	-0.060 [-0.097, -0.029]
2	R = 1; G = 0	0.072 [0.034, 0.116]	-0.018 [-0.050, 0.011]
2	R = 1; G = 1	0.049 [0.014, 0.088]	-0.014 [-0.046, 0.014]
2	ΔR; G = 0	0.030 [-0.021, 0.083]	0.089 [0.043, 0.144]
2	ΔR; G = 1	-0.020 [-0.074, 0.031]	0.046 [0.005, 0.090]

Note. ΔR denotes the high-minus-low reversibility difference in an indirect product. Percentile intervals use 5,000 respondent resamples within factorial cells, with seeds 20261016 and 20261017 for Studies 1 and 2. Intervals are unadjusted and characterize statistical indirect associations. Source: Authors' respondent-level analyses.

4.4 Final preference congruence

We estimate preference-congruence A × R interactions of − 0.009 [− 0.073, 0.055] in Study 1 and − 0.028 [− 0.056, − 0.001] in Study 2, contrary to H4's positive prediction. In Study 2, the unadjusted p = 0.040 becomes 0.080 after Holm correction; Study 1's adjusted p is 0.784. We find no support for H4.

We do not interpret Study 2's point estimate as evidence that undo access harms choice: it is an interaction, and its family-adjusted result is inconclusive. Differences in significance across psychological and behavioral outcomes also do not test differences between their effects. Our results show favorable psychological interactions alongside no demonstrated positive congruence interaction. We cannot attribute this pattern to recovery processes that we did not observe.

Note

Points show A × R contrasts with unadjusted 95% HC3 t intervals; Study 2 averages G equally. The acceptance panel shows an A × R contrast; H5 concerns G at A = 1, R = 0 (Table VI). Axes use different native units. Tables V and VI report Holm-adjusted p values for the seven-comparison family. Source: Authors' analyses.

4.5 Objective disclosure and recommendation acceptance

Under high autonomy and low reversibility, we observe a mean recommendation acceptance rate of 55.16% following consumer-priority disclosure and 46.25% following commercial-objective disclosure in

Study 2 (n = 160 per group). We estimate a difference of 0.0891 (HC3 SE = 0.0275), t(1272) = 3.235, p = 0.00125, 95% CI [0.0351, 0.1431]; Holm p = 0.00624 (Table VI). This 8.91 percentage-point difference supports H5. We do not test the disclosure contrast in Study 1 because G is fixed at 1.

After covariate adjustment, we estimate a difference of 0.0906 [0.0367, 0.1445], p = 0.00100. The respondent bootstrap interval is [0.0359, 0.1438], corroborating the positive association. We do not interpret this conditional contrast as evidence that disclosure matters only in this condition. The acceptance A × R × G interaction is − 0.1031 [− 0.2083, 0.0020], p = 0.0545, providing inconclusive evidence of three-way heterogeneity. We interpret H5 as an association with overall recommendation acceptance, without inferring purchase completion or improved discrimination.

4.6 Recommendation discrimination

In Study 2, we derive mean acceptance rates of 76.3% for matched and 34.0% for mismatched recommendations. Using 1,280 paired respondent summaries, we estimate a mean difference of 42.34 percentage points [39.93, 44.76]. This sample average does not imply successful discrimination by every respondent: accepting all recommendations maximizes uptake but produces a zero discrimination gap.

We estimate an average reversibility contrast on discrimination of − 0.62 percentage points [− 5.47, 4.22] and an average AR contrast of − 6.25 [− 15.94, 3.44]. We find no support for H6, although both intervals retain beneficial and adverse possibilities. These secondary intervals are unadjusted. Without event sequences, we cannot test whether acceptance preceded successful restoration.

4.7 Covariate sensitivity and hypothesis assessment

We assess sensitivity to age, gender, region and AI-use frequency (Table VIII). Encroachment interactions remain negative and continuation interactions positive in both datasets; we also report congruence and acceptance. We exclude unavailable residence type. These checks do not rule out unmeasured confounding or demographic differences in scale meaning.

Table VIII. Covariate sensitivity of the autonomy by reversibility contrasts

Study	AR contrast	Unadjusted	Covariate adjusted	Adjusted 95% CI
1	Encroachment	-0.423	-0.426	[-0.692, -0.160]
1	Continuation	0.369	0.393	[0.147, 0.640]
1	Congruence	-0.009	-0.012	[-0.078, 0.054]
1	Acceptance	0.063	0.061	[-0.089, 0.210]
2	Encroachment	-0.416	-0.422	[-0.602, -0.242]
2	Continuation	0.312	0.308	[0.136, 0.480]
2	Congruence	-0.028	-0.029	[-0.056, -0.002]
2	Acceptance	0.075	0.076	[0.023, 0.129]

Note. Covariates are age centered and scaled per ten years, gender, region and AI-use frequency. Study 2 AR contrasts average G equally. HC3 t intervals are not adjusted for multiplicity. This specification check does not test demographic moderation or measurement invariance. Source: Authors' respondent-level analyses.

Table VIII assesses covariate sensitivity; Tables V and VI address multiplicity within the seven-comparison family. All models retain 640 and 1,280 respondents as the analytical samples.

We summarize the hypothesis assessments in Table IX. We find composite-level associations consistent with H1 and H2, the predicted encroachment interaction and a compatible continuation pattern for H3, and a positive acceptance contrast for H5 after Holm correction. We find no support for H4's positive congruence prediction or H6's discrimination benefit. Our results therefore distinguish favorable evaluations and increased uptake from demonstrated improvements in choice protection.

Table IX. Hypotheses and empirical evidence

H	Relevant comparison	Observed pattern	Assessment
H1	A → offloading; offloading product	S1: 0.211; S2 average: 0.272; positive products	Consistent observed association; supervision costs unmeasured
H2	A → encroachment at low R	S1 A: 0.428; negative encroachment products	Consistent adverse association; permission-boundary mechanism unisolated
H3	AR for encroachment	S1: -0.423; S2 average: -0.416	Predicted encroachment interaction; causal mechanism unestablished
H4	AR for final congruence	S1: -0.009; S2: -0.028	No predicted positive pattern; Holm p = 0.784 and 0.080
H5	G → recommendation acceptance at A1 R0	55.16% vs 46.25%; Δ = 0.089; Holm p = 0.006	Supported for recommendation acceptance
H6	R → matched–mismatched gap	-0.62 pp [-5.47, 4.22]	No demonstrated improvement; pp = percentage points

Note. Hypothesis assessments follow the estimands and inference procedures in Sections 3.2–3.3. Holm adjustment covers the six A × R comparisons and the H5 disclosure contrast. Source: Authors' analyses.

5. Discussion

We find that higher autonomy is associated with greater offloading, while reversibility changes the autonomy contrasts in encroachment and continued delegation in the predicted directions. Consumer-priority disclosure also predicts higher recommendation acceptance under high autonomy and low reversibility. We find no demonstrated improvement in final congruence or discrimination. Our framework explains why consumer authority requires assessment of the delegated arrangement, the consumer's experience and the consequences of intervention.

We locate the delegation boundary at a particular action and decision stage (Baird and Maruping, 2021; Adam et al., 2024; Voyer et al., 2025). A consumer can authorize comparison while reserving seller selection, or authorize basket configuration while reserving payment. Our contribution is a framework for relating retained rights to outcomes that can diverge. We do not assume that the five measured perceptions exhaust consumer authority or that the current results fully validate the arrangement.

We interpret the encroachment and continuation interactions as consistent with reassurance from reversibility. We do not establish a corresponding improvement in final congruence or selective acceptance. This pattern does not show that restoration is ineffective or that effects on different scales differ statistically. It motivates separate assessment of mismatch detection, successful correction and final retention. A positive evaluation of correction access cannot establish successful correction.

Our H5 result links disclosed priorities to recommendation uptake. We do not equate this association with purchase conversion or objective alignment. Accurate disclosure can also justify refusing a revealed

commercial priority (Du and Xie, 2021; Nizette et al., 2025). We therefore evaluate reliance in relation to what consumers accept and the consequences of that choice (Lee and See, 2004).

5.1 Theoretical implications

We extend delegation theory by specifying how transferred execution rights coexist with retained intervention rights. Invocation, execution scope and restoration concern different decisions; delegating more work need not eliminate consumer authority. Our account yields a testable implication: holding the selected option and its quality constant, varying understood permission should change encroachment if the retained boundary is the mechanism. This comparison distinguishes our explanation from shared agency and autonomy-related aversion (Voyer et al., 2025; Rahman and Liang, 2026).

We also connect volitional endorsement with service governance. Self-determination does not require consumers to perform every operation personally (Ryan and Deci, 2020). We ask whether a delegated action remains within a scope the consumer endorses and can govern. Inspection, restoration and refusal perform different functions and can affect convenience differently. The amount of manual control alone is therefore insufficient to explain acceptance.

Within service-value research, we distinguish reassurance from choice protection to explain a possible divergence between evaluation and consequences (Grönroos, 2008; Ramaswamy and Ozcan, 2018; Chandra and Rahman, 2024). A favorable control rating can coexist with failed restoration, while an effective safeguard can impose a burden that lowers convenience. Our results motivate this distinction without establishing the unobserved restoration process. Measuring expectations before execution and recovery success afterward would provide a direct test.

We regard informed refusal as compatible with responsible service provision. Continuation and recommendation acceptance are relevant marketing outcomes, but their meaning depends on what is accepted and the consequences of the decision. Relating uptake to objective alignment, comprehension and final retention provides a stronger evaluation of consumer benefit than adoption alone.

5.2 Managerial implications

We recommend that providers specify the actions an agent may perform and obtain permission at the corresponding stage. Basket assembly, seller substitution and final submission can have different consequences, so permission to use AI need not cover all three. Our framework supports a limited permission menu for recoverable actions and a separate consumer decision when consequences cannot be restored. In Fig. 3, we translate this reasoning into a proposed authorization and correction process.

We recommend assessing restoration through recovery of the prior state, elapsed time and required steps. An undo control cannot guarantee recovery of price or stock. Task duration should include inspection and correction. Low correction rates require diagnosis because they can reflect missed mismatch, inaccessible intervention or endorsement of the retained choice.

We recommend comparing the disclosed ranking priority with the implemented rule and consumer understanding. Design teams should specify permission and recovery operations, commercial teams should specify optimization priorities, and the interface should communicate the arrangement. Separate logs of initial acceptance, restoration and final authorization would help distinguish informed decisions, provisional exploration and unwanted commitments. We connect responsible AI principles to these observable service decisions (Du and Xie, 2021).

We recommend pairing efficiency measures with a prespecified protection outcome. Claims of faster service without material quality deterioration require a justified non-inferiority margin and adequate precision; nonsignificance is insufficient. Purchase and firm-value claims additionally require transaction records. These recommendations define how providers can evaluate service performance beyond the outcomes observed here.

Note

We propose a process linking permission, execution and review to retention or restoration. The checks specify implementation responsibilities and costs; we do not evaluate an operating system. Source: Authors' conceptualization.

6. Limitations and future research

We analyse respondent-level composites and task summaries. Without event order or session identifiers, we cannot establish the sequence of acceptance and restoration or model shared task contexts. Our analytical records also do not verify the original recruitment, allocation, consent or ethics procedures. We therefore limit inference to associations within these datasets. Prospective replication should assess H5 using a prespecified outcome and testing plan, and observe purchase authorization and completed transactions separately.

We assess measurement quality using the accompanying summaries, which support internal consistency, convergent validity and HTMT point-estimate separation among the five included scales. We do not estimate a complete measurement model. Respondent-linked items, estimator details, global fit indices and bootstrap HTMT intervals are needed for fuller assessment. Configural, metric and scalar invariance tests would establish whether scores have comparable meanings across conditions. We cannot infer encroachment's distinction from unmeasured reactance or dissatisfaction from its separation from control and trust.

We recommend varying understood permission independently of product quality, measuring recovery expectations and observing restoration success to test the proposed mechanisms. Causal mediation also requires identification assumptions and sensitivity analysis (Imai et al., 2010). Prospective interaction power and preregistered outcomes and multiplicity families would strengthen confirmation. Our seven-comparison correction does not cover every outcome, and unmeasured boundary conditions cannot explain away estimates contrary to the hypotheses.

We see scope to extend the analysis to consequential purchases, repeated use and additional product categories. Preference congruence depends on the attributes and feasible alternatives, and consumers may revise their priorities during a task. Interviews and interaction logs could distinguish informed preference change from unrecognized overreach. Evaluating retained authority requires evidence about both consumers' understanding of the arrangement and the consequences of the final action.

Declarations

Author Contribution

The sole author was responsible for the conceptualization, study design, methodology, data analysis, interpretation of results, and manuscript preparation, including writing, reviewing, and editing. The author read and approved the final manuscript.

Data Availability

The datasets supporting the findings of this study are provided in the Supplementary Material accompanying this article.

Funding Declaration

This research was supported by Zhejiang Technical Institute of Economics.

References

1. Adam, M., C. Diebel, M. Goutier, and A. Benlian. 2024. Navigating autonomy and control in human-AI delegation: User responses to technology- versus user-invoked task allocation. *Decision Support Systems* 180:114193. https://doi.org/10.1016/j.dss.2024.114193
2. Ameen, N., A. Tarhini, A. Reppel, and A. Anand. 2021. Customer experiences in the age of artificial intelligence. *Computers in Human Behavior* 114. Article 106548. https://doi.org/10.1016/j.chb.2020.106548
3. Araujo, T. 2018. Living up to the chatbot hype: The influence of anthropomorphic design cues and communicative agency framing on conversational agent and company perceptions. *Computers in Human Behavior* 85:183–189. https://doi.org/10.1016/j.chb.2018.03.051
4. Ashfaq, M., J. Yun, S. Yu, and S. M. C. Loureiro. 2020. I, Chatbot: Modeling the determinants of users' satisfaction and continuance intention of AI-powered service agents. *Telematics and Informatics* 54:101473. https://doi.org/10.1016/j.tele.2020.101473
5. Baird, A., and L. M. Maruping. 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

6. Čaić, M., G. Odekerken-Schröder, and D. Mahr. 2018. Service robots: value co-creation and co-destruction in elderly care networks. *Journal of Service Management* 29(2):178–205. https://doi.org/10.1108/JOSM-07-2017-0179
7. Cameron, A. C., and D. L. Miller. 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources* 50(2):317–372. https://doi.org/10.3368/jhr.50.2.317
8. Candrian, C., and A. Scherer. 2022. Rise of the machines: Delegating decisions to autonomous AI. *Computers in Human Behavior* 134:107308. https://doi.org/10.1016/j.chb.2022.107308
9. Chandra, B., and Z. Rahman. 2024. Artificial intelligence and value co-creation: a review, conceptual framework and directions for future research. *Journal of Service Theory and Practice* 34(1):7–32. https://doi.org/10.1108/JSTP-03-2023-0097
10. Du, S., and C. Xie. 2021. Paradoxes of artificial intelligence in consumer markets: Ethical challenges and opportunities. *Journal of Business Research* 129:961–974. https://doi.org/10.1016/j.jbusres.2020.08.024
11. Fornell, C., and D. F. Larcker. 1981. Evaluating structural equation models with unobservable variables and measurement error. *Journal of Marketing Research* 18(1):39–50. https://doi.org/10.1177/002224378101800104
12. Frank, D.-A., L. F. Jacobsen, H. A. Søndergaard, and T. Otterbring. 2023. In companies we trust: consumer adoption of artificial intelligence services and the role of trust in companies and AI autonomy. *Information Technology & People* 36:155–173. https://doi.org/10.1108/ITP-09-2022-0721
13. Grönroos, C. 2008. Service logic revisited: who creates value? And who co-creates? *European Business Review* 20(4):298–314. https://doi.org/10.1108/09555340810886585
14. Gursoy, D., O. H. Chi, L. Lu, and R. Nunkoo. 2019. Consumers acceptance of artificially intelligent (AI) device use in service delivery. *International Journal of Information Management* 49:157–169. https://doi.org/10.1016/j.ijinfomgt.2019.03.008
15. Henseler, J., C. M. Ringle, and M. Sarstedt. 2015. A new criterion for assessing discriminant validity in variance-based structural equation modeling. *Journal of the Academy of Marketing Science* 43(1):115–135. https://doi.org/10.1007/s11747-014-0403-8
16. Holm, S. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* 6(2):65–70. https://www.jstor.org/stable/4615733
17. Huang, M.-H., and R. T. Rust. 2022. AI as customer. *Journal of Service Management* 33(2):210–220. https://doi.org/10.1108/JOSM-11-2021-0425
18. Imai, K., L. Keele, and D. Tingley. 2010. A general approach to causal mediation analysis. *Psychological Methods* 15(4):309–334. https://doi.org/10.1037/a0020761
19. Ivanova-Stenzel, R., and M. Tolksdorf. 2024. Measuring preferences for algorithms - How willing are people to cede control to algorithms? *Journal of Behavioral and Experimental Economics* 112:102270. https://doi.org/10.1016/j.socec.2024.102270
20. Jaakkola, E. 2020. Designing conceptual articles: four approaches. *AMS Review* 10:18–26. https://doi.org/10.1007/s13162-020-00161-0

21. Järvi, H., A.-K. Kähkönen, and H. Torvinen. 2018. When value co-creation fails: Reasons that lead to value co-destruction. *Scandinavian Journal of Management* 34(1):63–77. https://doi.org/10.1016/j.scaman.2018.01.002
22. Jensen, M. C., and W. H. Meckling. 1976. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics* 3(4):305–360. https://doi.org/10.1016/0304-405x(76)90026-x
23. Larivière, B., D. Bowen, T. W. Andreassen, W. Kunz, N. J. Sirianni, C. Voss, N. V. Wünderlich, and A. De Keyser. 2017. Service Encounter 2.0: An investigation into the roles of technology, employees and customers. *Journal of Business Research* 79:238–246. https://doi.org/10.1016/j.jbusres.2017.03.008
24. Lee, J. D., and K. A. See. 2004. Trust in automation: designing for appropriate reliance. *Human Factors* 46(1):50–80. https://doi.org/10.1518/hfes.46.1.50_30392
25. Logg, J. M., J. A. Minson, and D. A. Moore. 2019. Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes* 151:90–103. https://doi.org/10.1016/j.obhdp.2018.12.005
26. MacKinnon, J. G., and H. White. 1985. Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics* 29(3):305–325. https://doi.org/10.1016/0304-4076(85)90158-7
27. McLean, G., and K. Osei-Frimpong. 2019. Hey Alexa … examine the variables influencing the use of artificial intelligent in-home voice assistants. *Computers in Human Behavior* 99:28–37. https://doi.org/10.1016/j.chb.2019.05.009
28. Nizette, F., W. Hammedi, A. C. R. van Riel, and N. Steils. 2025. Why should I trust you? Influence of explanation design on consumer behavior in AI-based services. *Journal of Service Management* 36(1):50–74. https://doi.org/10.1108/JOSM-05-2024-0223
29. Pizzi, G., D. Scarpi, and E. Pantano. 2021. Artificial intelligence and the new forms of interaction: Who has the control when interacting with a chatbot? *Journal of Business Research* 129:878–890. https://doi.org/10.1016/j.jbusres.2020.11.006
30. Plé, L., and R. Chumpitaz Cáceres. 2010. Not always co-creation: introducing interactional co-destruction of value in service-dominant logic. *Journal of Services Marketing* 24(6):430–437. https://doi.org/10.1108/08876041011072546
31. Rahman, M. J., and H. Liang. 2026. When do you want more control over AI? The paradox of autonomy, task stakes, and distrust in AI aversion. *International Journal of Information Management* 89:103060. https://doi.org/10.1016/j.ijinfomgt.2026.103060
32. Ramaswamy, V., and K. Ozcan. 2018. What is co-creation? An interactional creation framework and its implications for value creation. *Journal of Business Research* 84:196–205. https://doi.org/10.1016/j.jbusres.2017.11.027
33. Risko, E. F., and S. J. Gilbert. 2016. Cognitive Offloading. *Trends in Cognitive Sciences* 20(9):676–688. https://doi.org/10.1016/j.tics.2016.07.002

34. Ryan, R. M., and E. L. Deci. 2020. Intrinsic and extrinsic motivation from a self-determination theory perspective: Definitions, theory, practices, and future directions. *Contemporary Educational Psychology* 61:101860. https://doi.org/10.1016/j.cedpsych.2020.101860
35. Saeed, K., and V. R. Prybutok. 2026. When utility meets ethics: a stakeholder perspective on agentic information systems delegation. *International Journal of Information Management* 86. Article 102976. https://doi.org/10.1016/j.ijinfomgt.2025.102976
36. Shin, D. 2021. The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI. *International Journal of Human-Computer Studies* 146. Article 102551. https://doi.org/10.1016/j.ijhcs.2020.102551
37. Vargo, S. L., P. P. Maglio, and M. A. Akaka. 2008. On value and value co-creation: A service systems and service logic perspective. *European Management Journal* 26(3):145–152. https://doi.org/10.1016/j.emj.2008.04.003
38. Voyer, B. G., M. Sangle-Ferriere, L. Sajtos, and B. Sung. 2025. The measurement of perceived shared agency in customer-artificial intelligence interactions. *Journal of Service Theory and Practice* 35(4):632–658. https://doi.org/10.1108/JSTP-08-2024-0266
39. Wang, E. J., P. R. Berthon, and J. John. 2026. The refined interaction value framework: advancing value co-creation in human-AI service systems in the AI-era. *Journal of Services Marketing*. Advance online publicationhttps://doi.org/10.1108/JSM-12-2025-0888
40. Wirtz, J., P. G. Patterson, W. H. Kunz, T. Gruber, V. N. Lu, S. Paluch, and A. Martins. 2018. Brave new world: service robots in the frontline. *Journal of Service Management* 29(5):907–931. https://doi.org/10.1108/josm-04-2018-0119

Figures

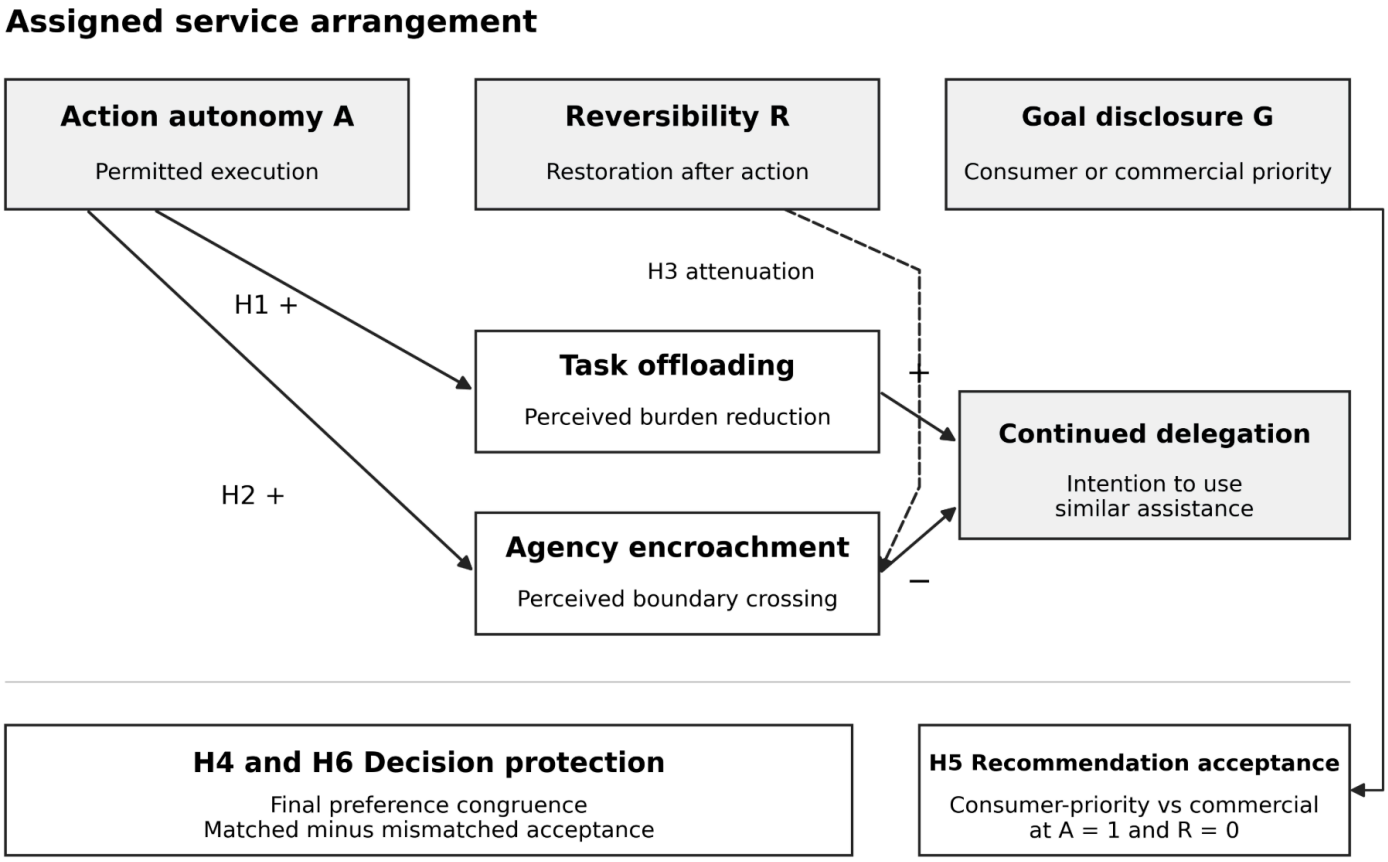

Figure 1

Permission arrangements and consumer outcomes

Note. A = action autonomy; R = reversibility; G = goal disclosure. H5 compares consumer-priority and commercial-objective disclosure on recommendation acceptance at A = 1, R = 0. Source: Authors' conceptualization.

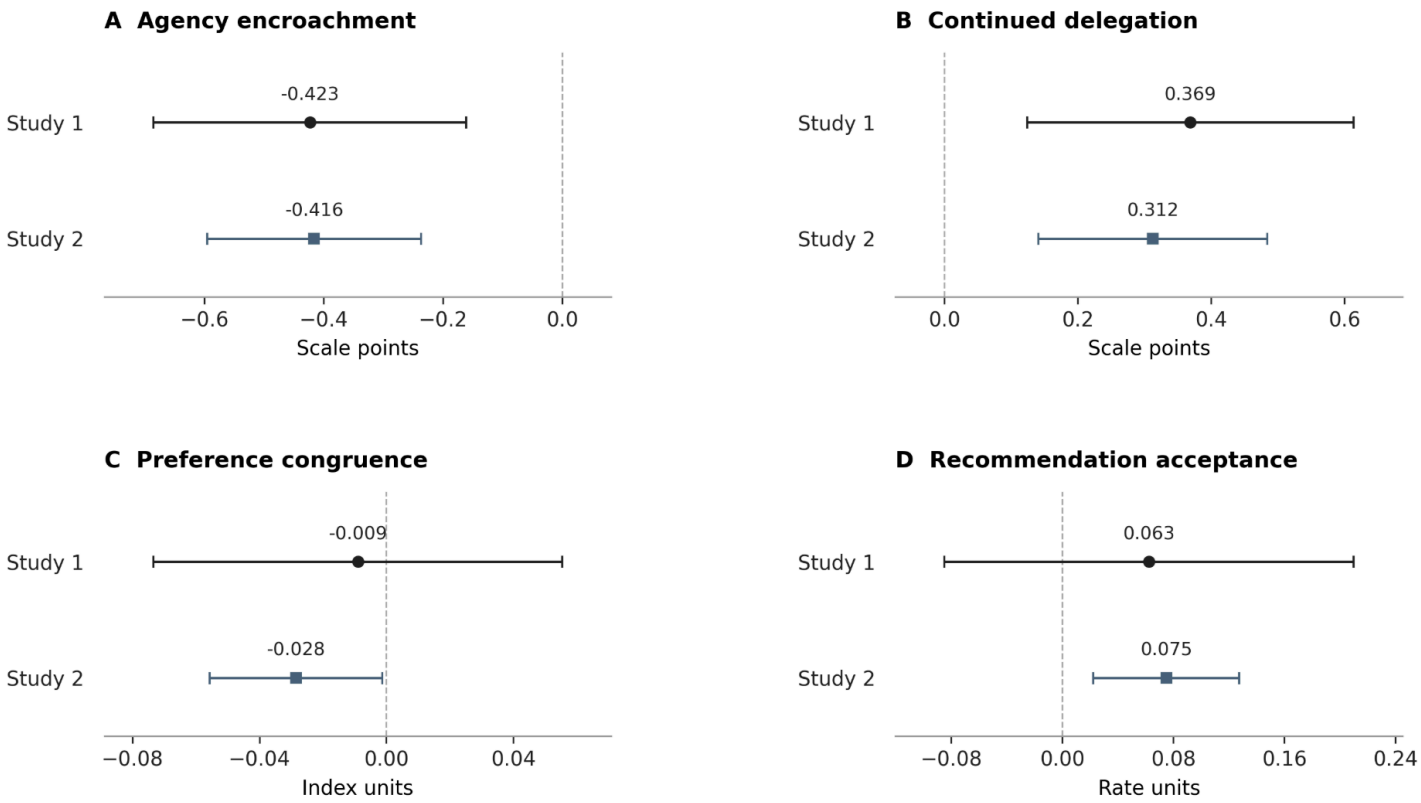

Figure 2

Autonomy by reversibility contrasts across outcomes

Note. Points show A × R contrasts with unadjusted 95% HC3 t intervals; Study 2 averages G equally. The acceptance panel shows an A × R contrast; H5 concerns G at A = 1, R = 0 (Table VI). Axes use different native units. Tables V and VI report Holm-adjusted p values for the seven-comparison family. Source: Authors' analyses.

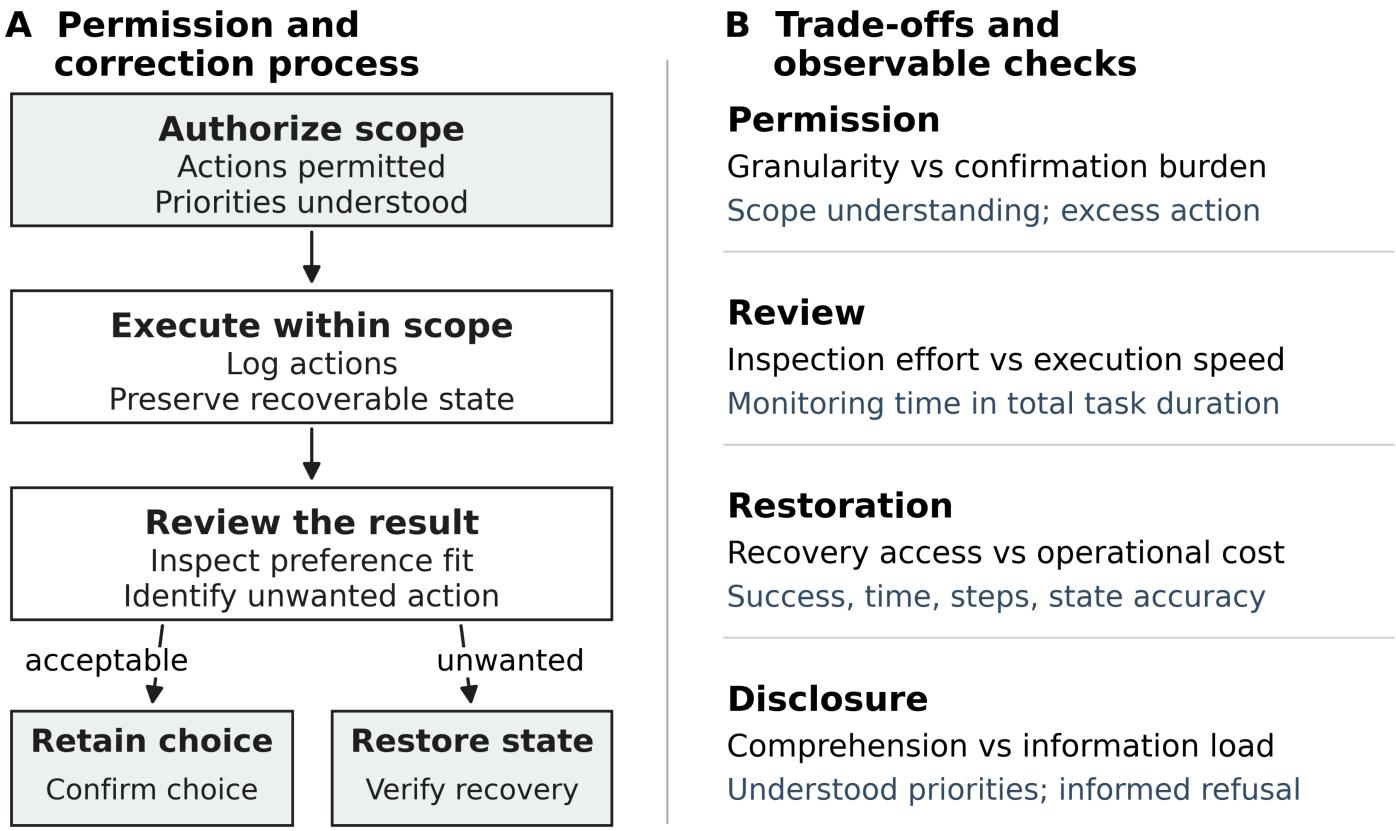

Figure 3

Authorization and correction in service implementation

Note. We propose a process linking permission, execution and review to retention or restoration. The checks specify implementation responsibilities and costs; we do not evaluate an operating system. Source: Authors' conceptualization.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Finaldata.xlsx