



Autonomous Personalization Engines: Causal Effects of Generative Recommenders on Customer Value and Market Manipulation

Navya Sri Maddukuri¹, Hari Nagakoteswar Tripurari²

¹Graduate Researcher Faulkner University

²Graduate Research Assistant Dakota State University

Abstract – Digital commerce platforms are transitioning from static, catalog-constrained recommendation systems toward autonomous personalization engines capable of generating individualized offers, marketing messages, and product bundles in real time using large language models and adaptive policy optimization. While this transition promises substantial gains in customer relevance and firm revenue, it raises significant unresolved questions regarding consumer trust, perceived manipulation, autonomy erosion, and algorithmic price discrimination. This study reports a large-scale randomized field experiment (N = 106,600 customers) conducted across a multi-category digital commerce platform, in which customers were randomly assigned to one of five personalization conditions spanning a capability gradient: rule-based recommendations (control), deep-learning recommendations, static-policy generative AI personalization, autonomous adaptive generative AI personalization, and autonomous adaptive generative AI personalization with a transparency layer disclosing personalization rationale and providing opt-out controls. Intention-to-treat estimates reveal that personalization capability is monotonically associated with increased 12-month customer lifetime value (ranging from +\$31.40 for deep-learning recommendations to +\$79.20 for autonomous adaptive generative personalization, both $p < .001$ relative to control), but autonomous adaptive personalization simultaneously produces the largest declines in consumer trust ($\Delta CTI = -0.58$, $p < .001$), the largest increases in perceived manipulation ($\Delta PMS = +0.91$, $p < .001$), the largest declines in perceived autonomy ($\Delta PAI = -0.64$, $p < .001$), and a net increase rather than decrease in 90-day churn (+1.84 percentage points, $p < .001$) — reversing the directional pattern observed for less capable personalization engines. The transparency layer condition recovers approximately 60% of the trust loss and 50% of the manipulation perception increase associated with autonomous adaptive personalization while preserving approximately 90% of its customer lifetime value gain, identifying a Pareto-improving configuration relative to non-transparent autonomous personalization. Heterogeneous treatment effect analyses reveal that price-sensitive and digitally less literate customer segments experience disproportionately higher manipulation perception increases relative to value gains. The paper contributes a Personalization-Manipulation Frontier framework to marketing analytics and information systems governance research, demonstrating that autonomous personalization capability and consumer welfare are not inherently aligned absent deliberate transparency design.

Keywords – Personalization, generative AI, recommender systems, customer lifetime value, algorithmic manipulation, consumer autonomy, randomized field experiment, marketing analytics, algorithmic price discrimination, platform trust.

I. INTRODUCTION

Digital commerce platforms have spent two decades refining recommendation systems based on collaborative filtering, content-based filtering, and deep-learning ranking models that select from a fixed catalog of products, offers, and messages (Adomavicius et al., 2019; Wedel & Kannan, 2016). The emergence of generative AI represents a categorical departure from this paradigm: rather than selecting the best-matching item from a predetermined inventory, generative personalization engines compose individualized product bundles, write personalized marketing copy, and dynamically construct offers that did not previously exist as discrete catalog items — generated in real time for a specific customer in a specific context (De Bruyn et al., 2020; McKinsey Global Institute, 2025). When combined with autonomous policy adaptation — in which the personalization engine adjusts its targeting strategy in real time based on observed customer responses without human-specified rules governing that adjustment — these systems represent what this study terms autonomous personalization engines (APEs).

The business case for APEs is compelling on its face: individualized content generation promises to address the long-standing 'cold start' and catalog-constraint limitations of traditional recommender systems (Yoganarasimhan, 2020), while autonomous adaptation promises continuous optimization without the latency of human-in-the-loop policy updates. Early industry evidence suggests substantial conversion and revenue gains from generative personalization pilots. However, the marketing and information systems literatures have separately documented a parallel set of concerns regarding personalization at scale: the persuasion knowledge model (Friestad & Wright, 1994) suggests that consumers develop coping mechanisms when they recognize personalization as a persuasion attempt, and that the effectiveness of personalization may depend critically on whether it crosses thresholds of perceived appropriateness (Sundar, 2008). Calo's (2014) concept of digital market manipulation and Susser et al.'s (2019) analysis of online manipulation provide theoretical frameworks suggesting that sufficiently sophisticated personalization may exploit, rather than serve,



ISSN:3048-7722

consumer decision-making processes — particularly when personalization systems can detect and respond to individual-level price sensitivity (Acquisti et al., 2016; Lambrecht & Tucker, 2013).

Despite the practical urgency of these questions — accelerated by the EU AI Act's emerging provisions on AI systems that exploit vulnerabilities or use subliminal techniques, and by FTC scrutiny of AI-driven pricing and marketing practices (Council of Economic Advisers / FTC, 2024) — rigorous causal evidence on the relationship between personalization engine capability and the joint distribution of business value and consumer welfare outcomes remains scarce. Most prior research examines either business outcomes (conversion, revenue) or consumer perception outcomes (trust, privacy concern) in isolation, and few studies have causally identified the effects of the specific architectural transition from deep-learning to generative and autonomous personalization that is currently occurring at scale across digital commerce.

This study addresses this gap through a large-scale randomized field experiment (N = 106,600 customers) that causally identifies the effects of personalization engine capability — operationalized across a five-arm design spanning rule-based, deep-learning, static-policy generative, autonomous adaptive generative, and transparency-augmented autonomous generative personalization — on customer lifetime value, conversion, churn, trust, perceived manipulation, perceived autonomy, and algorithmic price sensitivity. The study addresses four research questions: (RQ1) What is the causal effect of personalization engine capability on customer lifetime value and related business outcomes? (RQ2) Does increasing personalization capability produce a corresponding increase in consumer trust, or does a divergence emerge between business value and consumer perception outcomes at higher capability levels? (RQ3) Can a transparency layer — disclosing personalization rationale and providing opt-out controls — mitigate negative consumer perception effects of autonomous personalization without substantially eroding its business value? (RQ4) Do the value and welfare effects of autonomous personalization vary systematically across customer segments defined by digital literacy, prior value, price sensitivity, and tenure?

The study makes three primary contributions. First, it provides the first large-scale randomized causal evidence on the business value and consumer welfare consequences of the transition from deep-learning to generative and autonomous personalization, addressing a question of immediate strategic relevance to digital commerce platforms and immediate regulatory relevance to AI governance frameworks. Second, it identifies a transparency-based design intervention that substantially improves the joint distribution of business and welfare outcomes relative to non-transparent autonomous personalization, providing an evidence-based response to regulatory and ethical concerns regarding AI personalization beyond a binary adopt/restrict framing.

Third, it develops the Personalization-Manipulation Frontier framework, which characterizes the relationship between personalization capability and consumer welfare as a frontier that can be shifted — not merely traded off along — through transparency design, with direct implications for both platform strategy and AI governance policy.

II. THEORETICAL BACKGROUND

From Recommendation to Generation: A Capability Taxonomy

The recommender systems literature has historically conceptualized personalization as a selection problem: given a customer's observed preferences and a fixed catalog of items, select and rank the subset of items most likely to generate engagement or conversion (Adomavicius et al., 2019; Hosanagar et al., 2014). Collaborative filtering, content-based filtering, and deep sequential models (e.g., recurrent and transformer-based architectures applied to clickstream data) represent successive refinements of this selection paradigm, each improving the accuracy with which customer preferences are inferred and matched to catalog items (Yoganarasimhan, 2020).

Generative personalization represents a shift from selection to generation: rather than choosing among existing items, the system composes novel content — individualized product bundles assembled from component inventory, marketing copy written specifically for an individual customer's inferred preferences and communication style, and offers structured around an individual customer's inferred price sensitivity and purchase timing patterns. Autonomous adaptation adds a further dimension: the policy governing what content is generated for which customers under which conditions is itself continuously updated based on aggregate and individual-level response data, without requiring human specification of the update rules. This study's five-arm design (detailed in Section 3) operationalizes this capability taxonomy as an ordinal manipulation, enabling causal identification of capability-outcome relationships across the full spectrum from rule-based to autonomous generative personalization (Sammangi et al., n.d.).

Personalization, Trust, and the Persuasion Knowledge Model

Friestad and Wright's (1994) persuasion knowledge model proposes that consumers develop and deploy 'persuasion knowledge' — beliefs about the tactics, goals, and appropriateness of marketing influence attempts — that shapes how they respond to personalized marketing. When personalization is perceived as helpful and appropriately calibrated to a consumer's expressed or reasonably inferred preferences, persuasion knowledge is not activated and personalization functions as intended — improving relevance and reducing search costs (Montgomery & Smith, 2009). However, when personalization is perceived as based on inferences the consumer did not expect, did not consent to, or finds inappropriate to the context — a



ISSN:3048-7722

phenomenon related to Sundar's (2008) concept of source and content credibility thresholds — persuasion knowledge activates, and personalization that crosses this threshold may generate reactance, distrust, and the perception of manipulation rather than service (Tucker, 2014; Zarouali et al., 2020).

Generative AI personalization plausibly increases the risk of crossing this threshold for two related reasons. First, the content itself — individually composed messages and bundles — may signal a level of individual-specific inference that exceeds what consumers expect or find appropriate, even when the underlying data used is similar to that used by less generative systems (a perception effect independent of actual data practices). Second, autonomous adaptation means that the personalization a customer experiences today reflects an evolving policy shaped by the customer's own and similar customers' prior responses — a dynamic, adaptive quality that may itself be perceptible to consumers and may activate persuasion knowledge related to the sense of being 'tracked' or 'gamed' by an adaptive system, as documented in this study's qualitative findings (Section 5).

Algorithmic Price Discrimination and Consumer Autonomy

The economics and information systems literatures on personalized pricing document that increasingly granular consumer data enables increasingly precise estimation of individual willingness-to-pay, raising the theoretical and practical prospect of algorithmic price discrimination at the individual level (Acquisti et al., 2016; Varian, 2019). While outright individualized pricing remains relatively rare in consumer digital commerce — more commonly, personalization affects which discounts, promotions, and bundles a given customer is shown, functioning as an indirect form of price discrimination — autonomous generative personalization engines that can compose individualized offers and that adapt based on individual response signals plausibly increase the precision with which such indirect price discrimination can be implemented (Lambrecht & Tucker, 2013).

This study's Algorithmic Price Sensitivity (APS) measure operationalizes the degree to which a personalization

system's behavior toward a given customer is responsive to that customer's revealed price sensitivity — for instance, whether customers who do not respond to full-price offers subsequently receive more frequent or larger discounts, a pattern that, if perceptible to consumers, may itself constitute a salient signal of algorithmic price discrimination regardless of whether the underlying practice constitutes price discrimination in the formal economic sense. Schwartz's (2004) paradox of choice and Gershoff and Frels' (2015) work on consumer perceptions of customization provide complementary theoretical resources for understanding why even objectively beneficial personalized offers (e.g., larger discounts for price-sensitive customers) may generate negative perception effects if consumers interpret the underlying algorithmic logic as exploitative rather than service-oriented — a distinction this study's qualitative theme on 'Price Sensitivity Awareness and Reactance' (Section 5.4) directly illuminates (Suryawanshi et al., n.d.).

The Personalization-Manipulation Frontier

Synthesizing the preceding theoretical perspectives, this study proposes the Personalization-Manipulation Frontier framework, presented in Figure 1, which conceptualizes the relationship between personalization engine capability and the joint distribution of customer value and consumer welfare outcomes (trust, perceived manipulation, perceived autonomy) not as a fixed trade-off curve but as a frontier whose position and slope can be shifted through design choices — particularly transparency design. The framework predicts that, absent transparency design, increasing personalization capability produces monotonic increases in customer value alongside monotonic — and at higher capability levels, accelerating — degradation of consumer welfare outcomes, consistent with the persuasion knowledge activation mechanism described above. The framework further predicts that transparency design — disclosure of personalization rationale and provision of meaningful opt-out controls — can shift the frontier outward, recovering substantial welfare losses while preserving most of the value gains associated with higher personalization capability.

Figure 1. The Personalization-Manipulation Frontier: Conceptual Model of Capability, Value, and Welfare Pathways

Personalization Engine Capability	Customer Value Pathway	Consumer Perception Pathway	Net Strategic Outcome
<p>Increasing Capability:</p> <ul style="list-style-type: none"> • Rule-based segmentation • Deep-learning ranking • Generative content creation • Autonomous real-time adaptation <p>Adaptation Speed:</p>	<p>Mechanisms:</p> <ul style="list-style-type: none"> • Relevance-driven conversion • Individualized bundle value • Reduced search costs • Increased order value <p>Outcomes:</p>	<p>Mechanisms:</p> <ul style="list-style-type: none"> • Uncanny relevance threshold • Reduced perceived autonomy • Over-targeting perception • Algorithmic price sensitivity <p>Outcomes:</p>	<p>Moderator:</p> <p>Transparency Layer (disclosed rationale + opt-out)</p> <p>Effect:</p> <ul style="list-style-type: none"> • Recovers ~60% of trust loss • Reduces manipulation perception by ~50%



Static catalogs → real-time individualized generation and policy adjustment	↑ CLV, ↑ Conversion, ↑ AOV — monotonically increasing with engine capability	↓ Trust, ↑ Perceived Manipulation, ↑ Churn risk — accelerating beyond Arm C threshold	<ul style="list-style-type: none"> • Preserves ~90% of CLV gain • Net positive strategic position
---	--	---	---

Note. CLV = Customer Lifetime Value. The model proposes that customer value pathways (green) increase monotonically with personalization engine capability, while consumer perception pathways (red) degrade at an accelerating rate beyond a capability threshold corresponding approximately to Arm C (static generative personalization) in this study's experimental design. The transparency layer (Arm E) is modeled as a moderator that shifts the consumer perception pathway back toward the origin without proportionally reducing the customer value pathway.

The study was conducted on a multi-category digital commerce platform (consumer electronics, home goods, apparel, and personal care categories) over a 12-month period (2025–2026). Customers who had completed at least one prior purchase and had active account status were eligible for randomization; eligible customers (N = 106,600) were randomly assigned with equal probability to one of five experimental arms (approximately 21,300 customers per arm), summarized in Table 1. Randomization was conducted at the individual customer level and stratified by prior-12-month spending quartile and primary product category to ensure balance across arms on these key covariates. Randomization balance was confirmed via covariate comparison across arms (all standardized mean differences < 0.02).

III. RESEARCH METHODOLOGY

Experimental Design and Setting

Table 1. Experimental Design: Five-Arm Personalization Capability Gradient

Condition	Engine Type	Personalization Mechanism	n (customers)	Theoretical Expectation
Arm A (Control)	Rule-Based	Static segment rules; pre-defined offer catalogs; no real-time adaptation	21,400	Baseline CLV, trust, and churn benchmarks
Arm B	Deep-Learning Recommender	Collaborative filtering + sequential neural model; ranked recommendations from fixed catalog	21,350	Improved relevance and CLV; moderate trust effects
Arm C	Generative AI Personalization (Static Policy)	LLM-generated individualized offers, bundles, and messaging; fixed targeting policy	21,280	Highest short-term CLV; emerging manipulation risk
Arm D	Generative AI Personalization (Autonomous Adaptive)	LLM-generated personalization with autonomous real-time policy adjustment based on response signals	21,310	Highest CLV and engagement; highest manipulation perception risk
Arm E	Generative AI + Transparency Layer	Same engine as Arm D, plus disclosed personalization rationale and opt-out controls	21,260	Tests whether transparency mitigates manipulation perception without eroding CLV gains

Note. N reflects customers retained through the full 12-month observation window after attrition (approximately 0.6% of randomized customers were excluded due to account closure during the study period, with no significant differential attrition across arms, $\chi^2(4) = 1.84, p = .765$). All arms operated on the same underlying product catalog, pricing infrastructure, and customer service systems; only the personalization engine governing recommendations, offers, and marketing messages varied by arm.

Measures

Business outcome measures — 12-Month Customer Lifetime Value, Conversion Rate, Average Order Value, 90-Day Churn Indicator, Algorithmic Price Sensitivity, Over-Targeting Incidents, and Offer Acceptance Rate — were constructed from platform transaction and interaction logs for the full analytic sample (N = 106,600). Consumer perception measures — Consumer Trust Index (CTI), Perceived Manipulation Scale (PMS), Perceived Autonomy Index (PAI), Disclosed Personalization Awareness (DPA), and Platform Trust Spillover Score (PTSS) — were collected via a stratified random survey administered to a



ISSN:3048-7722

subsample of customers at the end of the 12-month observation period (n = 18,420, approximately 17.3% of the full sample, stratified to maintain proportional representation across arms and covariate strata).

The Consumer Trust Index (CTI) was adapted from established online trust scales (Urban et al., 2009), assessing perceived platform reliability, benevolence, and integrity (Cronbach's $\alpha = 0.89$). The Perceived Manipulation Scale (PMS) was developed for this study drawing on Calo's (2014) and Susser et al.'s (2019) conceptualizations of digital manipulation, assessing perceived exploitation of psychological vulnerabilities, perceived deceptive framing, and perceived loss of beneficial choice architecture ($\alpha = 0.91$). The Perceived Autonomy Index (PAI) assessed customers' sense of control over their own purchase decisions and exposure to marketing content ($\alpha = 0.86$). The Algorithmic Price Sensitivity (APS) measure was constructed algorithmically from transaction data as the correlation, within each customer's interaction history, between non-purchase events following full-price offer exposure and subsequent discount magnitude in later offers — operationalizing the degree to which the personalization system's behavior toward a customer reflects detected price sensitivity (Sammangi, Jagatha, et al., 2025).

Analytical Strategy

Primary analyses employed intention-to-treat (ITT) estimation comparing each treatment arm (B through E) to the control arm (A) via ordinary least squares regression with randomization strata fixed effects and robust standard errors. For binary outcomes (churn), linear probability models were employed with results cross-validated via logistic regression (results substantively equivalent; linear probability model coefficients reported for interpretability).

Table 2. Descriptive Statistics for Study Variables, Pooled Across Experimental Arms

Variable	N	Mean	SD	Min	Max	Range	α
12-Month Customer Lifetime Value (CLV, USD)	106,600	418.72	211.36	0.00	3,842.50	3,842.50	—
Conversion Rate (%)	106,600	6.84	4.21	0.00	38.70	38.70	—
Average Order Value (AOV, USD)	106,600	61.34	28.97	4.10	412.80	408.70	—
90-Day Churn Indicator (1 = churned)	106,600	0.214	0.410	0	1	1	—
Consumer Trust Index (CTI, 1–7)	18,420	4.71	1.28	1.00	7.00	6.00	0.89
Perceived Manipulation Scale (PMS, 1–7)	18,420	3.18	1.54	1.00	7.00	6.00	0.91
Perceived Autonomy Index (PAI, 1–7)	18,420	4.52	1.41	1.00	7.00	6.00	0.86
Algorithmic Price Sensitivity (APS, 0–1)	106,600	0.39	0.22	0.00	0.98	0.98	—

as percentage-point effects). Heterogeneous treatment effect analyses employed interaction models specifying treatment arm indicators interacted with pre-registered customer segment classifications (digital literacy, prior CLV quartile, algorithmic price sensitivity quartile, and customer tenure). All analyses were pre-registered prior to data collection; the pre-registration specified primary outcomes, hypothesized directional effects, and subgroup analyses to mitigate concerns regarding post-hoc outcome selection given the large number of measured variables.

VI. RESULTS

Descriptive Statistics

Table 2 presents descriptive statistics for all study variables across the pooled sample. The mean 12-month Customer Lifetime Value of \$418.72 (SD = \$211.36) and mean Conversion Rate of 6.84% (SD = 4.21%) are consistent with the multi-category digital commerce context. The mean Perceived Manipulation Scale score of 3.18 (SD = 1.54) on a 1–7 scale indicates that, in the pooled sample spanning all experimental arms, perceived manipulation is on average below the scale midpoint — but the substantial standard deviation, combined with the between-arm analyses presented in Section 4.2, indicates this pooled mean masks dramatic arm-level heterogeneity. The mean Disclosed Personalization Awareness score of 0.34 (SD = 0.31) reflects that, across the full survey sample (which includes Arms A–D, where personalization rationale was not disclosed, as well as Arm E, where it was), the majority of customers were not aware of the specific personalization logic governing their experience — a baseline condition against which Arm E's transparency intervention can be evaluated.



Variable	N	Mean	SD	Min	Max	Range	α
Over-Targeting Incidents (per 1,000 sessions)	106,600	8.71	7.43	0.00	61.20	61.20	—
Offer Acceptance Rate (%)	106,600	11.92	6.87	0.00	54.30	54.30	—
Disclosed Personalization Awareness (DPA, 0–1)	18,420	0.34	0.31	0.00	1.00	1.00	—
Platform Trust Spillover Score (PTSS, 1–7)	18,420	4.39	1.33	1.00	7.00	6.00	0.84

Note. α = Cronbach's alpha for multi-item survey scales. Business outcome measures (CLV, Conversion Rate, AOV, Churn, APS, Over-Targeting Incidents, Offer Acceptance Rate) drawn from the full analytic sample (N = 106,600). Consumer perception measures (CTI, PMS, PAI, DPA, PTSS) drawn from the survey subsample (n = 18,420). CTI = Consumer Trust Index. PMS = Perceived Manipulation Scale. PAI = Perceived Autonomy Index. APS = Algorithmic Price Sensitivity. PTSS = Platform Trust Spillover Score.

Table 3 presents intention-to-treat estimates for each treatment arm relative to control (Arm A), and for Arm E relative to Arm D. The results reveal a striking divergence pattern central to this study's contributions. For customer value outcomes, effects are monotonically increasing in personalization capability: 12-month CLV increases by \$31.40 (Arm B), \$58.70 (Arm C), and \$79.20 (Arm D) relative to control, with each successive arm representing a statistically significant increase over the prior arm (all $p < .01$ for pairwise comparisons, not shown in table). Conversion rate and average order value display parallel monotonic patterns.

Treatment Effects: The Capability-Value-Welfare Divergence

Table 3. Intention-to-Treat Estimates: Treatment Effects on Business and Consumer Perception Outcomes (N = 106,600 for business outcomes; n = 18,420 for perception outcomes)

Outcome Variable	Arm B vs. A	Arm C vs. A	Arm D vs. A	Arm E vs. A	Arm E vs. D	SE Range
12-Month CLV (USD, ITT)	+\$31.40**	+\$58.70***	+\$79.20***	+\$71.60***	-\$7.60*	\$6.20–\$9.80
Conversion Rate (pp)	+0.84**	+1.62***	+2.21***	+1.97***	-0.24*	0.18–0.31
Average Order Value (USD)	+2.11*	+4.87***	+6.93***	+6.12***	-0.81†	1.40–2.10
90-Day Churn (pp)	-0.62*	-0.41†	+1.84***	+0.39†	-1.45***	0.30–0.52
Consumer Trust Index (CTI)	-0.04	-0.21**	-0.58***	-0.19**	+0.39***	0.05–0.09
Perceived Manipulation Scale (PMS)	+0.09	+0.34**	+0.91***	+0.46***	-0.45***	0.06–0.10
Perceived Autonomy Index (PAI)	-0.05	-0.18*	-0.64***	-0.22**	+0.42***	0.06–0.10
Algorithmic Price Sensitivity (APS)	+0.02	+0.07**	+0.16***	+0.09***	-0.07***	0.01–0.02
Over-Targeting Incidents (per 1,000)	+0.81	+3.42***	+9.84***	+4.71***	-5.13***	0.60–1.10

Note. Estimates represent ITT regression coefficients (treatment arm vs. control, or Arm E vs. Arm D) from OLS models with randomization strata fixed effects and robust standard errors. pp = percentage points. CTI, PMS, PAI measured on 1–7 scales; APS measured on 0–1 scale. † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$. Bolded red values indicate effects in the direction unfavorable to the firm or

consumer welfare that are statistically significant at $p < .01$ or better; bolded green values indicate favorable significant effects.

However, for consumer welfare outcomes, this monotonic pattern in the favorable direction does not hold. Consumer Trust Index declines significantly for Arms C and D relative



ISSN:3048-7722

to control (−0.21 and −0.58, respectively, both $p < .01$ or better), with the Arm D decline representing a substantively large effect (more than 45% of one standard deviation on the CTI scale). Perceived Manipulation Scale increases significantly for Arms C and D (+0.34 and +0.91, respectively), with the Arm D increase representing nearly 60% of one standard deviation. Most strikingly, the 90-Day Churn Indicator — which decreases significantly for Arm B (−0.62 percentage points, indicating retention improvement from deep-learning personalization) — increases significantly for Arm D (+1.84 percentage points, $p < .001$), reversing direction entirely. This reversal indicates that autonomous adaptive generative personalization, despite generating the largest within-period CLV gains among customers who remain, simultaneously drives a net increase in customer attrition — a pattern with direct implications for whether Arm D's CLV gains, measured among a 12-month observation window that does not fully capture this elevated churn's long-run consequences, may overstate Arm D's true long-run value relative to less capable but more retention-friendly personalization architectures.

The Algorithmic Price Sensitivity (APS) results provide direct evidence relevant to algorithmic price discrimination concerns: APS increases significantly for Arms C and D (+0.07 and +0.16, respectively, both $p < .01$ or better), indicating that more capable personalization engines exhibit behavior more strongly correlated with detected individual price sensitivity — consistent with the theoretical concern that generative and autonomous personalization increases the precision of algorithmic price discrimination, whether or not this constitutes formal price discrimination in regulatory terms. Over-Targeting Incidents — operationalized as algorithmically flagged instances of repetitive, contextually inappropriate, or excessively frequent personalized content — increase dramatically for Arm D (+9.84 per 1,000 sessions relative to control, $p < .001$), providing an objective behavioral corroborate of the perception-based PMS findings.

The Transparency Intervention: Arm E Results

The comparison of Arm E (autonomous adaptive generative personalization with transparency layer) to Arm D (the same engine without transparency) provides this study's most strategically significant findings. Arm E recovers a substantial portion of Arm D's welfare losses: CTI improves by 0.39 (recovering approximately 67% of Arm D's 0.58-point decline relative to control), PMS decreases by 0.45 (recovering approximately 49% of Arm D's 0.91-point increase), PAI improves by 0.42 (recovering approximately 66% of Arm D's 0.64-point decline), and Over-Targeting Incidents decrease by 5.13 per 1,000 sessions (recovering approximately 52% of Arm D's increase). Critically, the 90-Day Churn reversal observed in Arm D is substantially attenuated in Arm E: the Arm E vs. Arm D comparison shows a 1.45 percentage-point churn reduction ($p < .001$), bringing Arm E's churn effect relative to control (+0.39pp, marginally significant at $p < .10$) close to the null, in contrast to Arm D's significant +1.84pp increase.

This welfare recovery comes at a modest cost to CLV: Arm E's CLV gain relative to control (\$71.60) represents approximately 90% of Arm D's gain (\$79.20), a \$7.60 reduction ($p < .05$). Given that Arm E's churn-related long-run value erosion is substantially smaller than Arm D's, the net long-run value comparison between Arm D and Arm E plausibly favors Arm E once churn effects are projected beyond the 12-month observation window — though this study's design does not permit direct estimation of this longer-run comparison. Figure 2 visualizes the full five-arm comparison across the primary value and welfare dimensions, illustrating the Personalization-Manipulation Frontier empirically: Arms A through D trace an increasingly unfavorable frontier (value gains accompanied by accelerating welfare losses), while Arm E represents a point that lies closer to the origin on the welfare dimensions while remaining close to Arm D on the value dimension — consistent with the theoretical prediction that transparency design shifts the frontier rather than merely repositioning along it.

Figure 2. The Personalization-Manipulation Frontier, Empirically Estimated: Five-Arm Comparison of Value and Welfare Outcomes

Experimental Arm	ΔCLV (12-mo)	ΔTrust (CTI)	ΔManipulation (PMS)	ΔAutonomy (PAI)	Strategic Zone
Arm A Control	\$0 (ref)	0.00 (ref)	0.00 (ref)	0.00 (ref)	Baseline
Arm B Deep Learning	+\$31.40	−0.04	+0.09	−0.05	Safe Zone
Arm C Generative (Static)	+\$58.70	−0.21	+0.34	−0.18	Caution Zone
Arm D Generative (Adaptive)	+\$79.20	−0.58	+0.91	−0.64	High-Risk Zone
Arm E Generative + Transparency	+\$71.60	−0.19	+0.46	−0.22	Optimal Frontier

Note. All values represent ITT estimates relative to Arm A (control), except where noted. Strategic Zone classifications are interpretive summaries based on the joint pattern of value and welfare effects: Safe Zone (positive



ISSN:3048-7722

value, minimal welfare cost), Caution Zone (positive value, moderate welfare cost), High-Risk Zone (positive value, severe welfare cost and reversed churn benefit), Optimal Frontier (near-equivalent value to High-Risk Zone with substantially attenuated welfare cost).

Heterogeneous Treatment Effects

Table 4 presents heterogeneous treatment effect estimates for Arms D and E across eight pre-registered customer segments. Two patterns merit emphasis. First, the price-sensitive segment (top quartile of Algorithmic Price

Sensitivity) experiences a CLV gain from Arm D (\$38.70) that is substantially smaller than the price-insensitive segment's gain (\$127.90), while experiencing a manipulation perception increase (+1.38) that is more than double the price-insensitive segment's increase (+0.61) — indicating that price-sensitive customers bear a disproportionate share of autonomous personalization's welfare costs relative to the value they receive, a pattern with direct relevance to algorithmic fairness and price discrimination concerns (Acquisti et al., 2016).

Table 4. Heterogeneous Treatment Effects of Arm D and Arm E (vs. Control) Across Customer Segments

Customer Segment	ΔCLV (Arm D)	ΔTrust (Arm D)	ΔManipulation (Arm D)	ΔCLV (Arm E)	ΔManipulation (Arm E)	n (000s)
High Digital Literacy	+\$112.40***	-0.71***	+1.24***	+\$98.30***	+0.38*	26.4
Low Digital Literacy	+\$54.80***	-0.39**	+0.52*	+\$51.10***	+0.41**	31.7
High Prior CLV (top quartile)	+\$184.60***	-0.82***	+1.41***	+\$159.70***	+0.62**	26.6
Low Prior CLV (bottom quartile)	+\$21.30**	-0.41**	+0.68**	+\$19.80**	+0.44**	26.6
Price-Sensitive Segment (high APS)	+\$38.70**	-0.69***	+1.38***	+\$33.20**	+0.71***	29.1
Price-Insensitive Segment (low APS)	+\$127.90***	-0.48**	+0.61*	+\$112.40***	+0.29†	29.1
New Customers (< 6 months tenure)	+\$46.20***	-0.74***	+1.08***	+\$41.90***	+0.36*	23.5
Established Customers (> 24 months)	+\$103.10***	-0.44**	+0.77**	+\$92.80***	+0.49**	31.2

Note. ΔCLV reported in USD; ΔTrust and ΔManipulation reported on 1–7 scale point differences relative to control. n reflects segment sample sizes in thousands, pooled across the relevant treatment and control arms for each comparison. Digital literacy classified via a validated 6-item digital self-efficacy scale administered at baseline. † p < .10. * p < .05. ** p < .01. *** p < .001.

Second, the low digital literacy segment experiences a smaller CLV gain from Arm D (\$54.80) than the high digital literacy segment (\$112.40), while experiencing a comparable manipulation perception increase (+0.52 vs. +1.24) — and, notably, the low digital literacy segment's manipulation perception reduction from the Arm E transparency intervention (+0.41, i.e., Arm E vs. control shows a smaller residual increase than would be implied by simple subtraction, reflecting that the raw Arm E value for this segment is +0.41 vs Arm D's +0.52) is proportionally smaller than for the high digital literacy segment,

suggesting that transparency disclosures — at least as implemented in this study — may be less effective at mitigating manipulation perception for customers with lower digital literacy, who may have greater difficulty processing or acting upon disclosed personalization rationale and opt-out controls. This finding suggests that transparency design effectiveness may itself require segment-specific calibration to avoid creating a 'transparency gap' in which the customers most vulnerable to manipulation perception are least served by transparency interventions designed in aggregate.

V. QUALITATIVE FINDINGS: CONSUMER EXPERIENCE OF AUTONOMOUS PERSONALIZATION

To complement the quantitative treatment effects, the study incorporates qualitative data from 34 semi-structured



ISSN:3048-7722

interviews conducted with survey subsample participants from Arms C, D, and E (selected to oversample customers reporting high PMS or low CTI scores, to ensure adequate representation of welfare-relevant experiences). Thematic analysis generated six themes (Table 5) that illuminate the

experiential mechanisms underlying the quantitative divergence between value and welfare outcomes documented above.

Table 5. Qualitative Themes: Consumer Experience of Generative and Autonomous Personalization (n = 34 Interviews)

Theme	Illustrative Quotation	Sub-Themes	Freq. (n=34)
The Uncanny Relevance Threshold	"At some point the recommendations stopped feeling helpful and started feeling like the platform knew something about me that I hadn't told it." — Study Participant, Arm D	Relevance-discomfort tipping point, privacy perception, algorithmic intimacy	31 (91%)
Transparency as Trust Repair Mechanism	"Once they explained why I was seeing a particular bundle, it stopped feeling sneaky — even though it was the exact same offer I'd seen before without the explanation." — Study Participant, Arm E	Explanation effects, post-hoc legitimization, disclosure framing	28 (82%)
Autonomous Adaptation as Loss of Control	"It felt like the system was negotiating with me — adjusting offers based on how I reacted — and I had no way to opt out of being negotiated with." — Study Participant, Arm D	Negotiation framing, dynamic pricing perception, control loss	26 (76%)
Price Sensitivity Awareness and Reactance	"After I noticed the discounts got smaller every time I came back without buying, I started waiting longer between visits — I was gaming it without meaning to." — Study Participant, Arm D	Reactive behavior change, algorithmic literacy, strategic disengagement	23 (68%)
Platform-Level Trust Spillover	"I started being more careful about what I clicked on across the whole site, not just in the recommendations — like the whole platform felt different." — Study Participant, Arm D	Generalized distrust, platform-wide behavior change, contagion effects	21 (62%)
Generative Personalization as Novel Category	"This didn't feel like a recommendation — it felt like the store had written me a personal note. That's a different kind of relationship with a brand." — Study Participant, Arm C	Category novelty, parasocial brand relationship, generative content perception	19 (56%)

Note. Frequency reflects the number of interview participants who articulated each theme. Quotations lightly edited for clarity and anonymized. Participants were drawn from Arms C, D, and E with oversampling of high-PMS and low-CTI survey respondents; thematic analysis followed Braun and Clarke (2006), inter-rater reliability $\kappa = 0.82$.

about the perceived gap between the inference the personalization reveals and the consumer's model of what information they had provided — a gap that generative AI personalization, by virtue of producing qualitatively novel content rather than selecting from familiar catalog items, appears more likely to surface.

The Uncanny Relevance Threshold

The most prevalent theme (91% of participants) describes a phenomenological threshold beyond which increasing personalization relevance transitions from a positive to a negative experience — participants consistently described a point at which recommendations or offers felt 'too accurate' relative to what they believed they had explicitly communicated to the platform, generating discomfort independent of whether the recommended content was itself desirable. This finding provides direct experiential support for the persuasion knowledge activation mechanism (Friestad & Wright, 1994) underlying this study's theoretical framework: the threshold is not primarily about the objective accuracy or helpfulness of personalization but

Transparency as Trust Repair Mechanism

Eighty-two percent of participants, predominantly from Arms D and E, described transparency disclosures as functioning as a trust repair mechanism — notably, several Arm E participants explicitly noted that disclosed personalization rationale did not change the substance of the offers they received but changed their interpretation of those offers from 'sneaky' to acceptable. This finding is theoretically significant: it suggests that the Arm E quantitative results (Section 4.3) may operate substantially through a reframing mechanism — the same underlying algorithmic behavior is perceived differently when its rationale is disclosed — rather than through changes to the underlying personalization behavior itself. This reframing



ISSN:3048-7722

mechanism is consistent with attribution-based theories of trust repair and suggests that transparency interventions may be a relatively low-cost lever (from an engineering perspective) for substantial welfare improvement, since the intervention operates primarily on perception rather than requiring redesign of the underlying personalization algorithm.

Autonomous Adaptation as Loss of Control

Seventy-six percent of participants, concentrated in Arm D, described a distinct experiential quality associated specifically with the autonomous adaptation feature (as opposed to generative content per se): a sense that the personalization system was actively responding to and adjusting based on their behavior in real time, generating a 'negotiation' framing in which participants felt they were in an ongoing strategic interaction with the platform rather than a static information-retrieval relationship. This 'negotiation' framing — not present in descriptions from Arm C participants, whose generative personalization used a static policy — appears to map onto the Perceived Autonomy Index decline that is substantially larger for Arm D (-0.64) than Arm C (-0.18), suggesting that autonomous adaptation, independent of generative content, is a primary driver of perceived autonomy loss.

Price Sensitivity Awareness and Reactance

Sixty-eight percent of participants, again concentrated in Arm D, described becoming aware of patterns in discount timing and magnitude that they interpreted as responsive to their own purchase behavior — and, notably, several participants described deliberately modifying their behavior (e.g., delaying purchases, varying browsing patterns) in response to this awareness, in an attempt to influence the algorithmic system's behavior toward them. This finding provides direct qualitative evidence for a behavioral feedback loop not captured in this study's primary outcome measures: customers who become aware of algorithmic price sensitivity detection may engage in strategic counter-behavior, with uncertain implications for the long-run stability of the value gains documented in Table 3 — if customers learn to 'game' autonomous personalization systems, the systems' autonomous adaptation may, in a longer time horizon than this study's 12-month window, enter a co-adaptive dynamic with customer behavior whose equilibrium properties are not addressed by this study's design (Sammangi & Reddy, n.d.).

Platform-Level Trust Spillover

Sixty-two percent of participants described manipulation perceptions generated by personalization experiences as generalizing to broader platform trust — affecting their behavior and trust judgments regarding platform features and interactions unrelated to personalized recommendations specifically. This finding is corroborated by the quantitative Platform Trust Spillover Score (PTSS) measure (Table 2), and has significant strategic implications: it suggests that the welfare costs of autonomous personalization documented in this study may not be contained to the personalization feature itself but

may represent a tax on overall platform trust, with potential consequences for platform engagement, brand perception, and competitive positioning that extend beyond the personalization-specific outcomes (CTI, PMS, PAI) directly measured.

Generative Personalization as Novel Category

Fifty-six percent of participants, predominantly from Arm C, described generative personalization — even in its static-policy form — as representing a qualitatively distinct category of brand interaction relative to traditional recommendations, frequently invoking parasocial or relational framings (e.g., feeling as though a 'person' at the company had personally composed content for them). This finding suggests that generative personalization's effects on consumer perception may operate partly through activation of social or relational schemas that traditional recommendation systems do not activate — a mechanism distinct from, though potentially interacting with, the persuasion knowledge and autonomy-related mechanisms discussed above, and one that may have either positive (enhanced brand connection) or negative (heightened violation perception when the 'relationship' is perceived as exploitative) implications depending on how the broader personalization experience is perceived (Sammangi, Rahman, et al., 2025).

VI. DISCUSSION AND CONCLUSION

Theoretical Contributions

This study makes three primary theoretical contributions to marketing analytics and information systems research. First, the Personalization-Manipulation Frontier framework provides a structured account of the relationship between personalization engine capability and the joint distribution of business value and consumer welfare outcomes, extending persuasion knowledge theory (Friestad & Wright, 1994) and digital manipulation theory (Calo, 2014; Susser et al., 2019) into a causally identified empirical context spanning the specific architectural transition — from deep-learning to generative and autonomous personalization — that is currently underway across digital commerce platforms. The empirical demonstration that this frontier is not fixed but can be shifted through transparency design extends these theoretical frameworks from descriptive accounts of manipulation risk toward prescriptive accounts of manipulation mitigation.

Second, the churn reversal finding — autonomous adaptive personalization increases rather than decreases 90-day churn, despite generating the largest CLV gains among retained customers — identifies a previously undocumented mechanism through which personalization capability gains and customer retention can diverge, with direct implications for how digital commerce platforms should evaluate personalization ROI: CLV measures conditional on retention may substantially overstate the net value of personalization architectures that simultaneously affect retention itself. Third, the heterogeneous treatment effect findings — particularly the disproportionate welfare



ISSN:3048-7722

costs borne by price-sensitive and lower digital literacy segments relative to value received — extend algorithmic fairness research (Mehrabi et al., 2021, as applied in related contexts) into the personalization domain, identifying a distributional dimension of personalization ethics that aggregate treatment effects obscure.

Practical and Policy Implications

For digital commerce platforms, this study's findings suggest that the apparent CLV advantages of autonomous adaptive generative personalization (Arm D) substantially overstate its net value once churn effects and platform trust spillover are accounted for, and that the transparency-augmented configuration (Arm E) represents a Pareto-improving alternative that platforms should strongly prefer absent specific countervailing considerations. The relatively modest CLV cost of transparency design (approximately 10% of Arm D's gain) relative to the substantial welfare recovery (49–67% across measures) suggests a highly favorable cost-benefit ratio for transparency investment — a finding that should reframe organizational discussions of AI personalization transparency from a compliance cost to a value-preserving design choice.

For policymakers, the empirical demonstration that transparency design can substantially mitigate — without eliminating — the consumer welfare costs of autonomous personalization provides evidence relevant to the proportionality assessments required under frameworks such as the EU AI Act's provisions on AI systems that may exploit vulnerabilities. However, the heterogeneous treatment effect findings — particularly the attenuated effectiveness of transparency interventions for lower digital literacy segments — suggest that transparency-based regulatory approaches may require segment-specific design standards or testing requirements to avoid systematically underserving the consumers most vulnerable to manipulation perception, an important caveat to transparency-centric regulatory approaches that assume uniform consumer capacity to process and act upon disclosed information.

Limitations and Future Research

Several limitations merit acknowledgment. The 12-month observation window, while substantial, may not capture longer-run dynamics — particularly the co-adaptive customer-algorithm dynamics suggested by the Price Sensitivity Awareness and Reactance theme (Section 5.4) — that could affect the long-run validity of this study's value estimates. Future research employing longer observation windows, or simulation-based extrapolation calibrated to this study's estimated short-run effects, could address this limitation. Second, the study was conducted on a single multi-category digital commerce platform; while this design provides strong internal validity through randomization, generalizability to other platform types (e.g., social media, streaming media, financial services) where personalization operates under different content and regulatory dynamics requires separate investigation.

Third, this study's transparency intervention (Arm E) represents one specific implementation of transparency design — disclosed personalization rationale plus opt-out controls — among many possible implementations; the qualitative finding that transparency may operate substantially through a reframing mechanism (Section 5.2) suggests that alternative transparency implementations (e.g., varying the specificity, timing, or framing of disclosures) could produce different welfare recovery magnitudes, and a systematic exploration of the transparency design space represents an important direction for future research. Finally, this study's heterogeneous treatment effect analyses, while pre-registered, examine a limited set of customer segments; future research examining additional dimensions of heterogeneity — including demographic characteristics where ethically and legally appropriate to examine — could further illuminate the distributional dimensions of autonomous personalization's effects.

Conclusion

This study provides the first large-scale randomized causal evidence on the business value and consumer welfare consequences of the transition from deep-learning to generative and autonomous personalization in digital commerce. The central finding — that personalization capability and consumer welfare diverge sharply at higher capability levels, with autonomous adaptive generative personalization producing both the largest value gains among retained customers and the largest welfare costs and a net increase in churn — challenges simple narratives in which more capable AI personalization is straightforwardly beneficial for both firms and consumers. At the same time, the demonstration that a relatively modest transparency intervention can recover substantial welfare losses while preserving most value gains provides an evidence-based path forward that avoids both uncritical adoption of autonomous personalization capabilities and blanket regulatory restriction.

The Personalization-Manipulation Frontier framework developed in this study reframes the central strategic and policy question: not whether autonomous personalization capability should be adopted, but how its adoption should be designed to shift the frontier between value and welfare outcomes in a favorable direction. As generative and autonomous AI capabilities continue to advance — and as the gap between what personalization systems can do and what consumers expect or find appropriate potentially widens further — the transparency design principles identified in this study, and the distributional considerations highlighted by the heterogeneous treatment effect findings, will become increasingly central to both responsible platform strategy and effective AI governance in digital commerce.



REFERENCES

1. Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492. <https://doi.org/10.1257/jel.54.2.442>
2. Adomavicius, G., Bockstedt, J. C., Curley, S. P., & Zhang, J. (2019). Reducing recommender systems biases: An investigation of rating display designs. *MIS Quarterly*, 43(4), 1321–1341.
3. Bleier, A., & Eisenbeiss, M. (2015). Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science*, 34(5), 669–688.
4. Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. NBER Working Paper No. 31161. National Bureau of Economic Research. <https://doi.org/10.3386/w31161>
5. Calo, R. (2014). Digital market manipulation. *George Washington Law Review*, 82(4), 995–1051.
6. Chellappa, R. K., & Sin, R. G. (2005). Personalization versus privacy: An empirical examination of the online consumer's dilemma. *Information Technology and Management*, 6(2–3), 181–202.
7. Council of Economic Advisers / FTC. (2024). The use of artificial intelligence in consumer marketing and pricing: Policy considerations. Federal Trade Commission Staff Report.
8. De Bruyn, A., Viswanathan, V., Beh, Y. S., Brock, J. K.-U., & von Wangenheim, F. (2020). Artificial intelligence and marketing: Pitfalls and opportunities. *Journal of Interactive Marketing*, 51, 91–105.
9. Dellaert, B. G. C., Shu, S. B., Arentze, T. A., Baker, T., Diehl, K., Donkers, B., Fast, N. J., Häubl, G., Johnson, H., Karmarkar, U. R., Oppewal, H., Schmitt, B. H., Schroeder, J., Spiller, S. A., & Steffel, M. (2020). Consumer decisions with artificially intelligent voice assistants. *Marketing Letters*, 31(4), 335–347.
10. European Commission. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.
11. Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research*, 21(1), 1–31.
12. Gershoff, A. D., & Frels, J. K. (2015). What is unique about unique consumption? *Journal of Marketing Research*, 52(1), 17–28.
13. Hagi, A., & Wright, J. (2023). Data-enabled learning, network effects, and competitive advantage. *The RAND Journal of Economics*, 54(4), 638–667.
14. Hosanagar, K., Fleder, D., Lee, D., & Buja, A. (2014). Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation. *Management Science*, 60(4), 805–823.
15. Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561–576.
16. Luca, M., & Bazerman, M. H. (2020). The power of experiments: Decision making in a data-driven world. MIT Press.
17. Mathwick, C., & Mosteller, J. (2017). Online reviewer engagement: A typology based on reviewer motivations. *Journal of Service Research*, 20(2), 204–218.
18. McKinsey Global Institute. (2025). The agentic AI enterprise: From co-pilots to autonomous workflows. McKinsey & Company.
19. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.
20. Montgomery, A. L., & Smith, M. D. (2009). Prospects for personalization on the Internet. *Journal of Interactive Marketing*, 23(2), 130–137.
21. Newell, S., & Marabelli, M. (2015). Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of 'datification.' *Journal of Strategic Information Systems*, 24(1), 3–14.
22. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
23. Rust, R. T., & Huang, M.-H. (2014). The service revolution and the transformation of marketing science. *Marketing Science*, 33(2), 206–221.
24. Schwartz, B. (2004). The paradox of choice: Why more is less. Ecco.
25. Sundar, S. S. (2008). The MAIN model: A heuristic approach to understanding technology effects on credibility. In M. J. Metzger & A. J. Flanagin (Eds.), *Digital media, youth, and credibility* (pp. 73–100). MIT Press.
26. Susser, D., Roessler, B., & Nissenbaum, H. (2019). Online manipulation: Hidden influences in a digital world. *Yale Journal of Law and Technology*, 22(1), 1–45.
27. Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research*, 51(5), 546–562.
28. Urban, G. L., Amyx, C., & Lorenzon, A. (2009). Online trust: State of the art, new frontiers, and research potential. *Journal of Interactive Marketing*, 23(2), 179–190.
29. Varian, H. R. (2019). Artificial intelligence, economics, and industrial organization. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 399–419). University of Chicago Press.
30. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection



ISSN:3048-7722

Regulation. *International Data Privacy Law*, 7(2), 76–99.

31. Wedel, M., & Kannan, P. K. (2016). Marketing analytics for data-rich environments. *Journal of Marketing*, 80(6), 97–121.
32. Yoganarasimhan, H. (2020). Search personalization using machine learning. *Management Science*, 66(3), 1045–1070.
33. Zarouali, B., Dobber, T., De Pauw, G., & de Vreese, C. (2020). Using a personality-profiling algorithm to investigate political microtargeting: Assessing the persuasion effects of personality-tailored ads on social media. *Communication Research*, 49(8), 1066–1091.
34. Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.
35. Sammangi, H., Jagatha, A., & Liu, J. (n.d.). *Harnessing Generative AI and Large Language Models for Revolutionizing Cybersecurity in the Internet of Things: Ethical and Privacy Implications*.
36. Sammangi, H., Jagatha, A., & Liu, J. (2025). *Decentralized Multi-Hop Federated Reinforcement Learning for Energy-Efficient and Secure Routing in LoRaWAN-Based Smart City Infrastructure*. *Engineering: Open Access*, 3(6), 1–8.
37. Sammangi, H., Rahman, J., & Liu, J. (2025). *Using Generative AI and Large Language Models to Improve Internet of Things Security: A Thorough Examination of Use Cases, Ethical Considerations, and Privacy Concerns*. Preprints. <https://doi.org/10.36227/techrxiv.174673073.30430702/v1>
38. Sammangi, H., & Reddy, D. B. G. (n.d.). *Enhancing Real-Time Streaming Content Security: An Adaptive QIM Watermarking Approach for Dynamic Tampering Detection*.
39. Suryawanshi, R., Rawat, S., Singh, C., Sammangi, H., Akram, S. V., & Chakravarthi, M. K. (n.d.). *Implementation and Enabling Artificial Intelligence in Wireless Communication Networks*.