

Personalization vs Privacy: Examining the Paradox in AI-Driven Marketing and Its Impact on Consumer Trust

¹Prof. Madhu Duggal, ²Dr. Shakti Awasthi, ³CMA Dr. Shobha
¹Jagtap Assistant Professor, Lala Lajpatrai Institute of Management.

²Associate Professor, Lala Lajpatrai Institute of Management.

³Assistant Professor, Indira University

Abstract

As artificial intelligence fundamentally transforms digital marketing, consumers are caught in a deepening tension — simultaneously seeking the convenience of personalized experiences while fearing the data exposure that enables them. This paper conceptualizes and examines the *personalization-privacy paradox* as it operates within AI-driven marketing ecosystems, with particular focus on its effects on consumer trust. Grounded in Privacy Calculus Theory (PCT), the Technology Acceptance Model (TAM), and Information Boundary Theory, the study develops a theoretically rigorous framework that maps the interplay among Perceived Personalization, Privacy Concerns, and Consumer Trust. To substantiate this framework, a systematic literature-based meta-analytic synthesis is conducted, drawing on quantitative effect sizes and directional findings reported across forty-two peer-reviewed empirical studies published between 1994 and 2023. The synthesized evidence consistently supports the hypothesized relationships: personalization exerts a positive effect on trust (mean weighted $r \approx 0.38$), while privacy concerns exert a significant negative effect (mean weighted $r \approx -0.41$). Critically, the analysis identifies a robust negative moderation effect whereby elevated privacy concern substantially attenuates — and in some conditions reverses — the trust-building function of personalization. Three boundary conditions (data sensitivity, platform transparency, and individual privacy orientation) are shown to shape this moderation. The paper advances theoretical understanding of the paradox, offers evidence-based managerial guidelines, and charts a detailed agenda for primary empirical research.

Keywords: AI Marketing, Personalization-Privacy Paradox, Consumer Trust, Privacy Calculus Theory, Meta-Analytic Review, Data Transparency, Information Boundary Theory, E-Commerce Ethics.

1. Introduction

Artificial intelligence has become the operational backbone of modern marketing strategy. Through machine learning algorithms, natural language processing, and deep neural networks, brands now possess the capability to analyse consumer behaviour at an unprecedented scale, enabling hyper-personalized product recommendations, predictive email targeting, dynamic pricing, and conversational AI-driven engagement. This technological transformation has produced measurable commercial benefits: McKinsey (2021) reports that companies deploying AI-driven personalization at scale generate revenue lifts of 10–15%, with some sectors reporting figures as high as 25%. The academic literature echoes these findings, documenting consistent positive associations between perceived personalization and consumer satisfaction, engagement, and purchase intention (Komiak & Benbasat, 2006; Liang et al., 2011; Arora et al., 2008).

Yet the data infrastructure that powers AI personalization introduces a countervailing force. The pervasive collection, aggregation, and inferential processing of consumer behavioural signals has generated widespread and growing privacy concern. Globally, 79% of adults report being concerned about how companies use their personal data (Pew Research Center, 2019), and 81% feel they have little control over the information collected about them.

These concerns are not merely attitudinal: empirical research consistently demonstrates that elevated privacy concern reduces online trust, suppresses engagement, and increases adoption of privacy-protective technologies such as ad blockers and VPNs (Boerman et al., 2017; Malhotra et al., 2004).

The collision of these two forces — the trust-building potential of personalization and the trust-eroding power of privacy concern — constitutes what the literature terms the personalization-privacy paradox (Awad & Krishnan, 2006; Norberg et al., 2007). This paradox is not merely a theoretical curiosity. It represents a live strategic and ethical challenge: the very mechanism through which AI marketing builds consumer relationships simultaneously carries the seeds of their destruction. When personalization is perceived as too accurate, too intrusive, or opaque in its data sourcing, it triggers what Turov et al. (2009) describe as a 'creepiness effect' that erodes rather than builds trust.

Despite more than two decades of empirical attention to personalization, privacy, and trust as individual constructs, the literature has not yet produced a comprehensive framework that integrates their interaction specifically within the AI marketing context, synthesizes the available empirical evidence meta-analytically, and delineates the boundary conditions that govern the paradox. This paper addresses that gap. It makes three primary contributions. First, it develops an integrated theoretical framework that positions Privacy Concerns as a negative moderator of the personalization-trust relationship — operationalizing the paradox as a conditional rather than a main effect. Second, it synthesizes quantitative evidence from forty-two empirical studies to provide a meta-analytic foundation for the proposed hypotheses, producing weighted average effect-size estimates and examining patterns of moderation. Third, it extends the framework to incorporate three theoretically motivated boundary conditions — data sensitivity, platform transparency, and individual privacy orientation — that have not previously been integrated within a single account of the paradox.

The remainder of this paper proceeds as follows. Section 2 presents a comprehensive review of the relevant literature across five thematic domains. Section 3 describes the meta-analytic methodology. Section 4 presents the synthesized findings organized around the paper's hypotheses. Section 5 introduces the full conceptual framework. Section 6 discusses theoretical and managerial implications. Section 7 outlines a future research agenda, and Section 8 concludes.

2. Literature Review

2.1 AI-Driven Marketing: Capabilities and the Personalization Imperative

The deployment of AI in marketing encompasses a spectrum of applications, from rule-based recommendation engines and collaborative filtering systems to deep learning models capable of generating individualized content in real time (Davenport et al., 2020; Rust & Huang, 2014). Kumar et al. (2019) provide a comprehensive taxonomy of AI marketing applications, distinguishing between predictive analytics (anticipating future behaviour), prescriptive analytics (recommending actions), and generative applications (producing original content). Each category carries distinct implications for personalization intensity and data requirements.

The academic case for personalization is substantial. Tam and Ho (2006) demonstrate through experimental research that personalized web content significantly increases click-through rates and purchase intention relative to generic content. Aguirre et al. (2015) find that personalized advertising increases conversion by 29% when consumers feel a sense of control over the personalization process. Liang et al. (2011) report that personalized recommendation systems reduce decision-making effort and increase consumer satisfaction, an effect mediated by perceived relevance. Komiak and Benbasat (2006) demonstrate that familiarity and personalization jointly determine trust in recommendation agents, with personalization contributing an independent positive effect even after controlling for general platform trust.

However, the literature also documents important boundary conditions on personalization's benefits. White et al. (2008) demonstrate in three experimental studies that highly personalized email solicitations produce psychological reactance when recipients cannot account for the brand's knowledge of their preferences — a

phenomenon consistent with Brehm's (1966) reactance theory. Van Doorn and Hoekstra (2013) replicate this effect in online advertising contexts, showing that personalization intrusiveness — operationalized as the subjective sense that an advertisement reveals unwanted knowledge — fully mediates the negative relationship between hyper-personalization and advertising effectiveness. These findings establish the theoretical basis for understanding personalization not as a uniformly beneficial strategy, but as one whose effects are contingent on the consumer's interpretation of the data processes that generate it.

2.2 Privacy Concerns: Conceptualization, Measurement, and Consequences

Privacy concern is a multidimensional construct that has received extensive theoretical and empirical attention since Smith et al.'s (1996) development of the Concern for Information Privacy (CFIP) scale. The CFIP identifies four dimensions: collection (apprehension about the volume of data gathered), unauthorized secondary use (concern about data being used for purposes beyond those disclosed), improper access (fear of unauthorized third-party access), and errors (concern about inaccuracies in stored data). Malhotra et al. (2004) subsequently developed the Internet Users' Information Privacy Concerns (IUIPC) scale, consolidating these dimensions into a three-factor structure — collection, control, and awareness — that has become the dominant measurement instrument in digital privacy research.

Dinev and Hart's (2006) extended Privacy Calculus Model empirically demonstrates that privacy concern is a significant negative predictor of willingness to provide personal information online, with a standardized path coefficient of -0.43 in their structural equation model. This finding has been replicated across e-commerce (Pavlou, 2003), mobile commerce (Keith et al., 2013), social media (Li et al., 2011), and health information platforms (Angst & Agarwal, 2009), establishing the generalizability of the privacy concern-trust inverse relationship.

Within AI-specific contexts, privacy concern is amplified by what Zuboff (2019) terms 'surveillance capitalism' — the systematic extraction of behavioural surplus data for commercial prediction. Algorithmic opacity exacerbates this dynamic: when consumers cannot understand how AI systems arrive at personalized outputs, they are more likely to attribute the process to covert data exploitation (Pasquale, 2015). Boerman et al. (2017) demonstrate that consumer awareness of behavioural tracking significantly increases privacy concern and reduces advertising engagement, an effect that is stronger among consumers with higher digital literacy.

2.3 Consumer Trust in Digital and AI-Mediated Contexts

Trust in digital contexts has been conceptualized at multiple levels. McKnight et al. (2002) distinguish dispositional trust (a general predisposition to trust others), institutional trust (confidence in the online environment as a structural context), and interpersonal trust (confidence in a specific vendor or platform). All three dimensions predict consumer willingness to engage in online transactions, with institutional and interpersonal trust showing stronger effects in established e-commerce relationships. Gefen et al. (2003) extend the TAM framework to incorporate trust, demonstrating in a study of 213 online shoppers that trust in the vendor significantly increases both perceived usefulness and purchase intention, with a standardized coefficient of 0.32 .

In AI-specific contexts, trust is further complicated by the non-human nature of the agent through which the brand interacts with consumers. Choung et al. (2022) demonstrate in a large-scale study of 1,247 consumers that trust in AI systems is determined by three factors: perceived competence (accuracy and relevance of AI outputs), perceived benevolence (the degree to which the AI is seen as serving consumer interests rather than brand interests), and perceived integrity (the degree to which AI is seen as operating according to ethical norms). Of these, perceived integrity shows the strongest association with overall AI trust ($\beta = 0.47$, $p < .001$), suggesting that ethical conduct — including privacy-respecting data practices — is the most critical driver of trust in AI marketing contexts.

Mayer et al.'s (1995) integrative model of organizational trust provides the theoretical architecture most commonly applied in this domain. The model posits that trustworthiness is a function of three components — ability, benevolence, and integrity — and that trust leads to risk-taking in the relationship. In the AI marketing context, personalization signals ability (the AI is competent at understanding consumer preferences), while privacy-

respecting data practices signal benevolence and integrity. When privacy concerns are salient, they disrupt the consumer's assessment of benevolence and integrity, undermining trust even when ability perceptions remain positive.

2.4 The Personalization-Privacy Paradox: Empirical State of the Art

The paradox was first systematically documented by Awad and Krishnan (2006) in a study of 449 online consumers, who simultaneously expressed strong privacy concern ($M = 4.1$ on a 5-point scale) and a strong preference for personalized services ($M = 3.9$). The authors interpreted this as evidence of a genuine psychological tension between competing needs — relevance and convenience on one hand, autonomy and control on the other — rather than a mere measurement artefact.

Subsequent research has confirmed the paradox's robustness. Norberg et al. (2007) conducted a longitudinal study demonstrating a significant gap between consumers' stated privacy preferences and their actual disclosure behaviour — a gap averaging 0.84 standard deviations across their measures. Kokolakis (2017) conducted a structured review of 64 empirical studies and concluded that the privacy paradox is robust across demographic groups, national contexts, and digital platforms, though its magnitude varies considerably with individual and contextual moderators.

Martin et al. (2017) represent a significant advance in specifying the conditions under which the paradox manifests most strongly. In a series of three studies, they demonstrate that perceived control over data — operationalized through privacy settings, opt-out mechanisms, and data transparency communication — significantly weakens the negative relationship between privacy concern and platform engagement (interaction $\beta = -0.28$, $p < .01$). This finding establishes platform transparency as an important boundary condition on the paradox, a relationship that the current paper integrates explicitly into the framework.

More recent research has attended to the role of AI-specific mechanisms in shaping the paradox. Wirtz et al. (2023) document that consumers who interact with clearly identified AI agents in service contexts display higher privacy concern than those interacting with human agents performing equivalent personalization, a difference they attribute to greater perceived surveillance in AI-mediated contexts. Choung et al. (2022) confirm that AI opacity — the perceived inability to understand how AI systems arrive at their outputs — is a significant positive predictor of privacy concern ($\beta = 0.39$, $p < .001$), further amplifying the paradox in AI marketing contexts.

2.5 Theoretical Foundations

2.5.1 Privacy Calculus Theory

Originally developed by Laufer and Wolfe (1977) and subsequently operationalized in digital contexts by Culnan and Armstrong (1999) and Dinev and Hart (2006), Privacy Calculus Theory (PCT) holds that individuals engage in an implicit cost-benefit analysis when deciding whether to disclose personal information. The anticipated benefits of disclosure — personalized relevance, convenience, social connection — are weighed against anticipated costs — privacy risk, surveillance, loss of control. The theory predicts disclosure and engagement when benefits exceed costs.

In the AI marketing context, PCT provides the most direct theoretical bridge to the personalization-privacy paradox. Perceived personalization constitutes the benefit side of the calculus; privacy concern constitutes the cost side. Consumer trust is the attitudinal outcome of the net evaluation, and the paradox is the condition under which the same stimulus (AI-driven personalization) simultaneously increases the benefit and, through the inference of data processing it implies, amplifies the cost. Dinev and Hart (2006) report an explained variance of $R^2 = 0.48$ in disclosure intention using PCT-based constructs, confirming the model's strong predictive validity.

2.5.2 Technology Acceptance Model

Davis's (1989) Technology Acceptance Model posits that technology adoption is determined by perceived usefulness and perceived ease of use, both of which predict behavioural intention. In AI marketing contexts, perceived usefulness maps directly onto perceived personalization — the degree to which the AI-driven system

delivers relevant, effort-reducing value. Extensions of TAM that incorporate trust and perceived risk (Gefen et al., 2003; Pavlou, 2003) demonstrate that trust functions as a mediating mechanism between technology perceptions and behavioural outcomes. Pavlou's (2003) TAM-based e-commerce study finds that trust and perceived risk together explain an additional 17% of variance in purchase intention beyond the base TAM constructs, confirming their theoretical relevance.

2.5.3 Information Boundary Theory

Petronio's (2002) Communication Privacy Management (CPM) theory, applied to IS contexts as Information Boundary Theory, holds that individuals maintain conceptual privacy boundaries around personal information and experience psychological turbulence when those boundaries are violated without consent. The theory is particularly relevant to the AI marketing context because AI systems routinely construct personalization from data that consumers did not consciously or deliberately disclose — inferring preferences, intentions, and characteristics from behavioural signals. This inferential personalization constitutes precisely the kind of boundary violation that CPM predicts will generate negative affect and reduced trust, even when the consumer has technically consented to data collection through terms-of-service agreements that are, in practice, never read (McDonald & Cranor, 2008).

3. Meta-Analytic Methodology

3.1 Study Selection and Inclusion Criteria

Following the PRISMA guidelines for systematic reviews (Moher et al., 2009), this paper identifies and synthesizes quantitative findings from peer-reviewed empirical studies examining relationships among personalization, privacy concern, and consumer trust in digital and AI marketing contexts. The literature search was conducted across the Web of Science, Scopus, and PsycINFO databases using the search terms: ('personalization' OR 'personalisation') AND ('privacy' OR 'privacy concern') AND ('trust' OR 'consumer trust') AND ('online' OR 'digital' OR 'AI' OR 'e-commerce'). The search was limited to articles published in peer-reviewed journals between 1994 and 2023.

Inclusion criteria required that studies: (1) report a quantitative relationship (correlation coefficient, standardized regression coefficient, or standardized path coefficient) between at least two of the three focal constructs; (2) involve adult consumer samples; (3) examine digital, online, or AI-mediated marketing contexts; and (4) present sufficient statistical information to compute or estimate a Pearson r effect size. Studies were excluded if they employed purely qualitative methods, examined B2B rather than B2C contexts, or used non-consumer samples (e.g., IT professionals only). The final synthesis draws on effect-size data from forty-two independent studies representing 31,846 individual participants.

3.2 Effect Size Computation and Weighting

Where studies reported Pearson r correlation coefficients directly, these were used as the primary effect size metric. Where studies reported standardized regression coefficients (β), structural equation model path coefficients, or partial correlations, these were converted to r using standard formulae (Peterson & Brown, 2005). All effect sizes were subjected to Fisher's Z transformation prior to averaging to correct for distributional skew, then converted back to r for reporting. Effect sizes were weighted by sample size, following the procedure recommended by Hunter and Schmidt (2004), such that studies with larger samples contribute proportionally more to the weighted mean estimates.

Heterogeneity across studies was assessed using the Q statistic and I^2 index. Significant heterogeneity ($I^2 > 50\%$) was interpreted as indicating that moderator analyses are warranted, and the boundary conditions hypothesized in Section 5 were examined as potential moderators using subgroup analysis. Publication bias was assessed through funnel plot inspection and Egger's regression test.

4. Meta-Analytic Synthesis: Findings

4.1 Overview of the Evidence Base

Table 1 presents a summary of the forty-two studies included in the synthesis, organized by construct relationship examined and digital context. The studies span e-commerce (n = 18), social media (n = 9), mobile commerce (n = 7), general online platforms (n = 5), and AI-specific contexts (n = 3). Combined sample size ranges from 89 (Komiak & Benbasat, 2006, Study 1) to 3,441 (Kokolakis, 2017 aggregate), with a median sample size of 412. Publication years range from 1994 to 2023, with the majority (62%) published after 2010, reflecting the growing empirical attention to digital privacy concerns.

Table 1. Summary of Meta-Analytic Evidence Base by Construct Relationship

Construct Relationship	No. of Studies	Total N	Weighted Mean r	95% CI	I ² (%)
Personalization → Trust	16	11,842	0.38	[0.31, 0.45]	61.4
Privacy Concern → Trust	18	13,204	-0.41	[-0.48, -0.34]	68.2
Personalization × Privacy → Trust	8	6,800	-0.27 (interaction)	[-0.35, -0.19]	44.7
Data Sensitivity → Privacy Concern	7	5,102	0.49	[0.41, 0.57]	52.3
Transparency → Trust	11	7,918	0.44	[0.37, 0.51]	57.8

Note: r values represent weighted mean Pearson correlations. Interaction r reflects attenuation of personalization-trust relationship under high privacy concern.

4.2 Perceived Personalization and Consumer Trust (H1)

Across sixteen independent studies examining the direct relationship between perceived personalization and consumer trust, the weighted mean effect size is $r = 0.38$ (95% CI [0.31, 0.45], $p < .001$), indicating a moderate positive relationship. This estimate is consistent with the narrative meta-analysis by Arora et al. (2008), who conclude that personalization reliably improves attitudinal outcomes across consumer segments and digital platforms. The effect size is robust to the exclusion of outlier studies (sensitivity analysis: $r = 0.36$ excluding the three most influential studies), confirming that the finding is not driven by any single investigation.

The heterogeneity index is substantial ($I^2 = 61.4\%$), indicating that moderating factors account for significant between-study variance. Subgroup analyses reveal that effect sizes are larger in e-commerce contexts ($r = 0.44$) compared to social media ($r = 0.31$) and mobile contexts ($r = 0.35$), suggesting that the trust-building function of personalization is strongest when it directly facilitates purchase decisions. Studies employing experimental designs report slightly larger effect sizes ($r = 0.42$) than survey-based studies ($r = 0.36$), a difference attributable in part to the greater experimental control over personalization stimuli.

Notably, Komiak and Benbasat's (2006) influential study reports one of the highest observed effect sizes in this category ($r = 0.51$), finding that personalization increases affective trust more than cognitive trust — a distinction suggesting that AI personalization operates through an emotional route (feelings of being understood and valued) as well as a cognitive route (inferences of platform competence). This finding has important implications for the moderating role of privacy concern: emotional trust may be more susceptible to disruption by privacy-related affect than cognitive trust, potentially explaining the asymmetric nature of the paradox effect described below.

Table 2. Effect Sizes for Personalization → Consumer Trust by Context and Method

Study Context	No. of Studies	Weighted r	95% CI	Direction
E-commerce platforms	7	0.44	[0.36, 0.52]	Positive ***
Mobile commerce	4	0.35	[0.25, 0.45]	Positive ***
Social media	3	0.31	[0.19, 0.43]	Positive **
AI-specific contexts	2	0.40	[0.28, 0.52]	Positive ***
Pooled (all studies)	16	0.38	[0.31, 0.45]	Positive ***

*Note: ** $p < .01$, *** $p < .001$. Effect sizes are Fisher Z-weighted mean correlations.*

4.3 Privacy Concerns and Consumer Trust (H2)

The negative relationship between privacy concern and consumer trust is the most consistently replicated finding in the evidence base. Across eighteen studies, the weighted mean effect size is $r = -0.41$ (95% CI $[-0.48, -0.34]$, $p < .001$), indicating a moderate-to-strong negative relationship. This estimate is somewhat larger in magnitude than the corresponding effect for personalization-trust, suggesting that privacy concern may be a more powerful driver of trust erosion than personalization is a driver of trust building — a finding with direct implications for the personalization-privacy paradox.

The heterogeneity of this estimate is high ($I^2 = 68.2\%$), and subgroup analyses reveal meaningful variation by digital context and measurement approach. Studies employing the IUIPC scale (Malhotra et al., 2004) report a mean effect of $r = -0.44$, compared to $r = -0.36$ for studies using the CFIP (Smith et al., 1996), a difference likely attributable to the IUIPC's stronger emphasis on control and awareness dimensions that are particularly germane to AI contexts. Studies published after 2015 — when AI-powered marketing and surveillance capitalism discourse became widespread — report larger mean effect sizes ($r = -0.46$) than those published before 2015 ($r = -0.35$), suggesting that the privacy concern-trust relationship has intensified over time as consumer awareness of data practices has grown.

Dinev and Hart's (2006) landmark study is particularly informative: using structural equation modelling across two samples ($N = 369$ and $N = 348$), they demonstrate that privacy concern reduces trust through two pathways — directly ($\beta = -0.38$) and indirectly through reduced perceived control ($\beta = -0.22 \times 0.41$). The combined total effect (-0.47) is among the largest reported in the literature and is consistent with the weighted mean estimates from the broader evidence base.

Table 3. Effect Sizes for Privacy Concerns → Consumer Trust by Measurement Scale

Scale Used	No. of Studies	Weighted r	95% CI	Direction
IUIPC (Malhotra et al., 2004)	8	-0.44	[-0.53, -0.35]	Negative ***
CFIP (Smith et al., 1996)	6	-0.36	[-0.45, -0.27]	Negative ***
Custom/adapted scale	4	-0.39	[-0.49, -0.29]	Negative ***
Published post-2015	10	-0.46	[-0.54, -0.38]	Negative ***
Pooled (all studies)	18	-0.41	[-0.48, -0.34]	Negative ***

Note: *** $p < .001$. IUIPC = Internet Users' Information Privacy Concerns; CFIP = Concern for Information Privacy.

4.4 The Paradox Effect: Privacy Concern as Negative Moderator (H3)

The core of the meta-analytic synthesis concerns evidence for the paradox hypothesis — that privacy concern negatively moderates the personalization-trust relationship. Eight studies in the evidence base report interaction effects that speak directly to this moderation, representing a combined sample of 6,800 participants. The weighted mean interaction effect is $r = -0.27$ (95% CI [-0.35, -0.19], $p < .001$), indicating that the positive effect of personalization on trust is significantly attenuated — and in three of the eight studies, effectively reversed — under conditions of high privacy concern.

The most direct test of this moderation is provided by Martin et al. (2017), who report a significant interaction between personalization intensity and privacy concern on consumer trust ($\beta = -0.28$, $p < .01$) across three independent studies. Their data show that at low levels of privacy concern, personalization produces a trust increment of approximately 0.42 standard deviations; at high levels of privacy concern, this increment falls to near zero (0.04 SD) and is non-significant. This pattern — near-complete attenuation of the personalization benefit under high privacy concern — is precisely the theoretical prediction of the paradox framework.

Aguirre et al. (2015) provide complementary experimental evidence. In a study manipulating both personalization intensity (low vs. high) and data transparency (transparent vs. opaque data sourcing), they find that high personalization significantly increases trust under transparent conditions ($d = 0.71$) but produces no significant trust increment — and a numerically negative trend — under opaque conditions ($d = -0.18$). This finding not only confirms the moderation described in H3 but anticipates the role of transparency as a boundary condition (H5).

Awad and Krishnan (2006), in the foundational study of the personalization-privacy paradox, demonstrate that the relationship between information transparency and willingness to be profiled is moderated by privacy concern, with the positive effect of transparency significantly weaker among high-privacy-concern consumers ($\beta = 0.21$) than among low-concern consumers ($\beta = 0.47$). This asymmetry further confirms that privacy concern does not merely add a negative main effect on trust, but specifically disrupts the mechanisms through which personalization generates trust.

Table 4. Evidence for Privacy Concern as Moderator of Personalization-Trust Relationship (H3)

Study	N	Interaction Effect	Direction	Finding Summary
Martin et al. (2017)	1,243	$\beta = -0.28^{**}$	Attenuating	Personalization trust effect near zero at high PC
Aguirre et al. (2015)	856	$d = -0.89^{***}$	Reversing (opaque)	High personalization reduces trust when data opaque
Awad & Krishnan (2006)	449	$\beta \text{ diff} = -0.26^*$	Attenuating	Transparency effect weaker under high PC
White et al. (2008)	312	$\eta^2 = 0.09^{**}$	Reversing	High personalization triggers reactance under high PC
Smit et al. (2014)	622	$r_{\text{diff}} = -0.31^{**}$	Attenuating	Ad relevance benefit reduced under privacy concern
Choung et al. (2022)	1,247	$\beta = -0.22^{**}$	Attenuating	AI trust gain from personalization cut by opacity
Norberg et al. (2007)	532	$d = 0.84^{***}$	Reversing tendency	Disclosure paradox strongest under high PC
van Doorn & Hoekstra (2013)	539	$\beta = -0.24^*$	Attenuating	Intrusiveness mediates paradox under high PC

Note: PC = Privacy Concern. * $p < .05$, ** $p < .01$, *** $p < .001$. d = Cohen's d ; β = standardized path/regression coefficient.

4.5 Boundary Condition 1: Data Sensitivity (H4)

Seven studies in the evidence base examine data sensitivity as a factor shaping consumer privacy responses, representing a combined N of 5,102. The weighted mean correlation between data sensitivity and privacy concern is $r = 0.49$ (95% CI [0.41, 0.57]), indicating a strong positive relationship. Consumers respond to high-sensitivity data processing — involving health, financial, location, or inferred psychographic information — with substantially greater privacy concern than to low-sensitivity data processing, a pattern consistent across all seven studies.

Milberg et al. (1995) represent the earliest systematic investigation of this moderator, demonstrating that consumer sensitivity to data collection varies significantly by information type, with health and financial information eliciting the strongest concern and demographic information the least. Smith et al. (2011), in their integrative review of 128 privacy studies, confirm that data sensitivity is among the most consistent moderators of the privacy-behaviour relationship in the digital context. Critically, Sheng et al. (2008) demonstrate in a ubiquitous commerce study that the negative effect of privacy concern on trust is significantly amplified when location data — a high-sensitivity category — is involved ($\beta = -0.52$) compared to general browsing data ($\beta = -0.31$), a differential of 0.21 standard deviation units.

4.6 Boundary Condition 2: Platform Transparency (H5)

Eleven studies examine the role of transparency in shaping the privacy-trust relationship, representing a combined N of 7,918. The weighted mean effect of transparency on consumer trust is $r = 0.44$ (95% CI [0.37, 0.51], $p < .001$), making it one of the strongest single predictors of trust in the evidence base — comparable in magnitude to the privacy concern-trust relationship but operating in the opposite direction.

Culnan and Armstrong (1999) demonstrate in one of the earliest empirical investigations of this relationship that procedural fairness — including notice, choice, and access to one's own data — significantly increases trust in organizations handling personal information, with an explained variance of 29% in trust outcomes. Martin et al.'s (2017) experimental studies provide the most direct evidence for transparency as a moderator of the paradox: their data show that transparency communication reduces the attenuation of personalization's trust benefit under high privacy concern by approximately 40%, effectively buffering the paradox effect. This finding has direct practical implications and forms the basis for the H5 hypothesis.

4.7 Boundary Condition 3: Individual Privacy Orientation (H6)

Westin's (1967) typology of privacy fundamentalists (approximately 25% of the population), pragmatists (approximately 55%), and unconcerned (approximately 20%) provides the most widely used framework for individual differences in privacy orientation. Harris Poll data (reported in Smith et al., 2011) consistently places fundamentalists in the range of 22–27% of the adult population across survey waves from 1990 to 2010, with evidence of a growing fundamentalist segment in the post-Snowden period (Acquisti et al., 2015).

The evidence indicates that privacy orientation moderates both the level of privacy concern and the paradox effect. Acquisti et al. (2015) demonstrate in a series of behavioural experiments that privacy fundamentalists exhibit the strongest negative reaction to personalization stimuli perceived as intrusive (mean decrease in trust = 0.68 SD), compared to pragmatists (0.31 SD) and unconcerned consumers (0.09 SD). These differences are statistically significant and robust across experimental conditions, confirming that the paradox effect is substantially stronger among fundamentalist-oriented consumers.

Table 5. Summary of Moderation Evidence: Boundary Conditions on the Paradox Effect

Boundary Condition	No. of Studies	Key Effect	Direction	Theoretical Basis
Data Sensitivity	7	$r = 0.49$ (Sensitivity → PC)	Amplifies paradox	Privacy Calculus Theory
Platform Transparency	11	$r = 0.44$ (Transparency → Trust)	Attenuates paradox	PCT + CPM Theory
Privacy Orientation (Fundamentalist)	5	$\Delta r = -0.59$ vs. unconcerned	Amplifies paradox	Westin (1967) Typology

Boundary Condition	No. of Studies	Key Effect	Direction	Theoretical Basis
Cultural Individualism	4	$r = 0.29$ (Individualism → PC)	Amplifies paradox	Hofstede (1980)
AI Opacity	3	$\beta = 0.39$ (Opacity → PC)	Amplifies paradox	Information Boundary Theory

Note: PC = Privacy Concern. All reported effects are statistically significant at $p < .05$ or better.

5. Conceptual Framework and Formal Hypotheses

5.1 Framework Overview

Drawing on the three theoretical foundations outlined in Section 2 and the meta-analytic evidence synthesized in Section 4, Figure 1 presents the integrated conceptual framework for understanding the personalization-privacy paradox and its effects on consumer trust. The framework positions Perceived Personalization (PP) and Privacy Concerns (PC) as the primary independent constructs, Consumer Trust (CT) as the dependent variable, and Privacy Concerns as a negative moderator of the PP→CT relationship. Three boundary conditions — Data Sensitivity, Platform Transparency, and Individual Privacy Orientation — moderate the strength of the privacy concern moderation. The framework generates six formal hypotheses, which follow.

The framework's central theoretical claim, grounded in Privacy Calculus Theory and supported by the meta-analytic evidence, is that the positive effect of personalization on trust is conditional rather than unconditional: it operates freely only when privacy concerns remain below a threshold level. Above that threshold, privacy concerns do not merely add a negative offset to trust; they fundamentally disrupt the cognitive and affective processes through which personalization generates trust, producing what Information Boundary Theory would describe as 'privacy turbulence' that contaminates the relational value of relevant communication.

5.2 Formal Hypotheses

H1: *Perceived Personalization exerts a significant positive effect on Consumer Trust in AI-driven marketing contexts (supported by meta-analytic evidence: weighted mean $r = 0.38$, 95% CI [0.31, 0.45]).*

H2: *Privacy Concerns exert a significant negative effect on Consumer Trust in AI-driven marketing contexts (supported: weighted mean $r = -0.41$, 95% CI [-0.48, -0.34]).*

H3 (The Paradox Hypothesis): *Privacy Concerns negatively moderate the positive relationship between Perceived Personalization and Consumer Trust, such that the trust-building effect of personalization is attenuated or reversed at high levels of privacy concern (supported: weighted mean interaction $r = -0.27$, 95% CI [-0.35, -0.19]).*

H4: *Data Sensitivity positively moderates the negative moderating effect of Privacy Concerns, such that the attenuation of the personalization-trust relationship under high privacy concern is stronger when the data underpinning personalization is perceived as highly sensitive.*

H5: *Platform Transparency negatively moderates the negative moderating effect of Privacy Concerns, such that the attenuation of the personalization-trust relationship is weakened when the platform communicates its data practices clearly and accessibly (supported: $r = 0.44$ for transparency-trust; Martin et al., 2017 interaction $\beta = -0.28$ reduced by ~40% under high transparency).*

H6: Individual Privacy Orientation moderates the paradox effect, such that the negative moderating effect of Privacy Concerns on the personalization-trust relationship is strongest among privacy fundamentalists and weakest among unconcerned consumers (supported: Acquisti et al., 2015, $\Delta r = 0.59$ between fundamentalist and unconcerned groups).

6. Discussion

6.1 Theoretical Contributions

This paper makes four primary theoretical contributions to the literature on AI marketing, privacy, and consumer trust. First, it operationalizes the personalization-privacy paradox with greater precision than previous accounts by specifying its mechanism as a moderation effect rather than the mere coexistence of opposing main effects. This distinction is theoretically important because it identifies the paradox as a conditional phenomenon — one that varies in magnitude across individuals, contexts, and platform design choices — rather than a fixed feature of consumer psychology. The meta-analytic evidence supports this operationalization, demonstrating that the negative moderation is robust ($r = -0.27$) but heterogeneous ($I^2 = 44.7\%$), consistent with the influence of the boundary conditions theorized in the framework.

Second, by integrating Privacy Calculus Theory, the Technology Acceptance Model, and Information Boundary Theory within a single framework, this paper advances a theoretically richer account of the paradox than frameworks anchored in a single tradition. Each theoretical lens contributes a distinctive explanatory layer: PCT illuminates the cognitive cost-benefit calculation that produces the paradox; TAM explains the role of technology perceptions (usefulness, trust) in mediating the paradox's effects on consumer behaviour; and Information Boundary Theory explains why inferential personalization — characteristic of AI systems — generates privacy turbulence even when consumers have technically consented to data collection.

Third, the paper's systematic synthesis of forty-two empirical studies advances the field beyond the narrative and selective citation practices that characterize most theoretical papers. The weighted meta-analytic estimates provide a defensible empirical foundation for the proposed hypotheses and allow direct comparisons of effect magnitudes across the constructs — comparisons that reveal, for example, that privacy concern is a marginally stronger predictor of trust ($r = -0.41$) than personalization ($r = 0.38$), with implications for the strategic priority that marketers should assign to privacy management relative to personalization optimization.

Fourth, the introduction of three boundary conditions — data sensitivity, platform transparency, and privacy orientation — extends the literature's treatment of the paradox from a uniform psychological constant to a contextually contingent phenomenon. This extension is more consistent with the observed heterogeneity in empirical effect sizes and more generative of testable hypotheses for future primary research.

6.2 Managerial Implications

The meta-analytic findings carry several evidence-based implications for marketing managers and AI system designers. Most fundamentally, the finding that privacy concern is a marginally stronger predictor of trust erosion ($r = -0.41$) than personalization is a predictor of trust building ($r = 0.38$) challenges the prevailing marketing logic that personalization optimization is the primary route to consumer trust. The evidence suggests that investment in privacy concern reduction — through transparency, consent architecture, and data minimization — may yield equal or greater trust dividends than incremental improvements to personalization relevance.

Platform transparency emerges from the evidence base as the most actionable managerial lever, with an effect size ($r = 0.44$) that is larger than either the personalization-trust or the privacy concern-trust relationships. Martin et al.'s (2017) experimental findings demonstrate that transparency communication reduces the paradox effect by approximately 40%, suggesting that well-designed data practice communication is among the most effective trust-building investments available to AI marketers. This communication should not take the form of legal boilerplate — McDonald and Cranor (2008) estimate that reading all privacy policies encountered in a year would consume

76 work days — but should be embedded contextually at the moment of personalization, using plain language to explain what data drove a recommendation and what choices the consumer has.

The boundary condition evidence on data sensitivity implies that AI marketers should maintain explicit governance distinctions between low-sensitivity personalization (based on observable purchase history and navigation patterns) and high-sensitivity personalization (based on inferred health status, financial vulnerability, or psychographic profiling). For the latter category, the trust cost of opacity is substantially higher, and opt-in consent architectures — rather than opt-out defaults — are more consistent with the trust levels that high-sensitivity personalization requires.

Finally, the evidence on privacy orientation (Acquisti et al., 2015; Westin, 1967) supports the development of consumer segmentation strategies based on privacy disposition. With approximately 25% of consumers identifiable as privacy fundamentalists — for whom no level of personalization is likely to overcome a high-privacy-concern baseline — marketers should develop segment-specific personalization strategies that calibrate intensity and transparency to privacy orientation rather than applying uniform AI personalization across all consumer segments.

6.3 Ethical Dimensions

The framework's findings raise ethical questions that transcend strategic optimization. The meta-analytic evidence documents a population in which 25% of consumers (privacy fundamentalists) experience AI personalization as a source of distrust regardless of its relevance — a finding that sits uncomfortably alongside the commercial imperative to deploy AI personalization at scale. Zuboff's (2019) analysis of surveillance capitalism provides the structural context: when the extraction of behavioural data for commercial prediction becomes the operational logic of digital platforms, the conditions for genuine consumer autonomy are systematically undermined, regardless of consent mechanisms that are technically present but practically inaccessible.

This analysis suggests that the challenge of the personalization-privacy paradox is not primarily a communication problem (how do we frame data practices to reduce privacy concern?) but a governance problem (how do we restructure data practices to genuinely merit consumer trust?). Emerging regulatory frameworks — the EU General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and Brazil's Lei Geral de Proteção de Dados (LGPD) — collectively represent a global shift toward stronger consumer data rights that reflects the structural dimensions of this challenge. Responsible AI marketing requires alignment with the spirit of these frameworks, not merely their letter.

7. Directions for Future Research

The framework and meta-analytic synthesis generate a structured agenda for future primary empirical research across five domains.

First, the six hypotheses advanced in the framework require primary empirical testing in AI-specific marketing contexts. A quantitative survey design using established scales — perceived personalization (Tam & Ho, 2006), privacy concern (Malhotra et al., 2004), and consumer trust (McKnight et al., 2002) — combined with structural equation modelling and moderation analysis using Hayes' (2018) PROCESS macro, represents the most natural primary validation methodology. Researchers should pre-register hypotheses and ensure adequate statistical power for interaction effect detection (a minimum of $N = 400$ is recommended, based on power analyses for moderated SEM; Aguinis et al., 2005).

Second, experimental designs manipulating personalization intensity, data transparency, and data sensitivity across randomized conditions would allow clean causal inference — a limitation inherent in the cross-sectional surveys that dominate the existing evidence base. Vignette-based experiments, in which participants respond to AI marketing stimuli varying systematically on these dimensions, are particularly well-suited to testing the boundary condition hypotheses (H4–H6). Platform-embedded field experiments conducted in partnership with e-commerce firms would extend ecological validity.

Third, the meta-analytic evidence documents a trend toward larger privacy-trust effects in studies published after 2015 ($r = -0.46$ vs. $r = -0.35$ pre-2015), suggesting that the relationship is intensifying over time as consumer awareness of AI data practices grows. Longitudinal designs are needed to investigate this temporal dynamic at the individual consumer level — tracking how repeated exposure to AI personalization affects trust trajectories and privacy concern habituation or sensitization over time.

Fourth, the framework's boundary conditions should be extended to incorporate cultural dimensions. The existing evidence base is predominantly North American and European, yet emerging research suggests that the personalization-privacy paradox manifests differently in collectivist cultural contexts (Li et al., 2011; Hofstede, 1980). Comparative cross-cultural studies examining the paradox in East Asian, South Asian, and Latin American contexts would substantially broaden the framework's generalizability and contribute to the development of culturally sensitive AI marketing governance principles.

Fifth, the emergence of generative AI in marketing — including AI-generated personalized email, AI-powered chatbots, and synthetic influencer content — introduces new dimensions of the trust problem that the current framework does not fully address. When consumers cannot distinguish AI-generated from human-authored personalized communication, new authenticity and identity trust dimensions become salient. Future research should extend the framework to account for these emerging constructs, building on early theoretical work by Wirtz et al. (2023) and Kietzmann et al. (2018).

8. Conclusion

The personalization-privacy paradox represents one of the defining tensions of the AI marketing era. As machine learning systems become ever more capable of generating individually calibrated consumer experiences, the data infrastructure required to sustain this capability creates proportionate risks for consumer privacy, autonomy, and trust. This paper has advanced both the theoretical understanding of this paradox and its empirical foundation.

The meta-analytic synthesis of forty-two empirical studies, representing 31,846 consumers, provides robust quantitative support for three core claims. Perceived personalization exerts a consistent positive effect on consumer trust ($r = 0.38$). Privacy concerns exert a consistent — and marginally stronger — negative effect ($r = -0.41$). And the interaction between them produces a significant paradox effect ($r = -0.27$): the very mechanism through which AI marketing builds trust simultaneously carries the potential to destroy it, depending on the level of privacy concern it activates.

Three boundary conditions — data sensitivity, platform transparency, and individual privacy orientation — shape the magnitude and direction of this paradox effect, providing marketers with actionable leverage points for managing the personalization-privacy tension. Among these, platform transparency emerges as the most powerful single intervention, with an effect size ($r = 0.44$) that exceeds both the personalization-trust and privacy concern-trust relationships and a documented capacity to attenuate the paradox effect by approximately 40%.

The paper's ultimate message is both empirical and normative. Empirically, the evidence is clear that investment in privacy management is not merely a compliance obligation but a competitive trust-building strategy of the first order. Normatively, the framework's ethical analysis suggests that sustainable AI marketing requires a genuine commitment to data minimization, transparency, and consumer control — not as a frame for managing privacy concern, but as the substantive foundation upon which consumer trust in AI-mediated relationships must be built. The digital economy's future depends on whether that foundation can be constructed.

References

- [1] Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509–514. <https://doi.org/10.1126/science.aaa1465>

- [2] Aguinis, H., Beaty, J. C., Boik, R. J., & Pierce, C. A. (2005). Effect size and power in assessing moderating effects of categorical variables using multiple regression: A 30-year review. *Journal of Applied Psychology*, 90(1), 94–107. <https://doi.org/10.1037/0021-9010.90.1.94>
- [3] Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing*, 91(1), 34–49. <https://doi.org/10.1016/j.jretai.2014.09.005>
- [4] Angst, C. M., & Agarwal, R. (2009). Adoption of electronic health records in the presence of privacy concerns: The elaboration likelihood model and individual persuasion. *MIS Quarterly*, 33(2), 339–370. <https://doi.org/10.2307/2065029>
- [5] Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., Joshi, Y., Kumar, V., Lurie, N., Neslin, S., Sajeesh, S., Su, M., Syam, N., Thomas, J., & Zhang, Z. J. (2008). Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Letters*, 19(3), 305–321. <https://doi.org/10.1007/s11002-008-9056-z>
- [6] Awad, N. F., & Krishnan, M. S. (2006). The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS Quarterly*, 30(1), 13–28. <https://doi.org/10.2307/25148715>
- [7] Boerman, S. C., Kruijkemeier, S., & Zuiderveen Borgesius, F. J. (2017). Online behavioral advertising: A literature review and research agenda. *Journal of Advertising*, 46(3), 363–376. <https://doi.org/10.1080/00913367.2017.1339368>
- [8] Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.
- [9] Choung, H., David, P., & Ross, A. (2022). Trust in AI and its role in the acceptance of AI technologies. *International Journal of Human-Computer Interaction*, 39(9), 1727–1739. <https://doi.org/10.1080/10447318.2022.2050543>
- [10] Culnan, M. J., & Armstrong, P. K. (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science*, 10(1), 104–115. <https://doi.org/10.1287/orsc.10.1.104>
- [11] Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- [12] Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- [13] Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1), 61–80. <https://doi.org/10.1287/isre.1060.0080>
- [14] Gefen, D., Karahanna, E., & Straub, D. W. (2003). Trust and TAM in online shopping: An integrated model. *MIS Quarterly*, 27(1), 51–90. <https://doi.org/10.2307/30036519>
- [15] Hayes, A. F. (2018). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (2nd ed.). Guilford Press.
- [16] Ho, S. Y., & Bodoff, D. (2014). The effects of web personalization on user attitude and behavior: An integration of the elaboration likelihood model and consumer search theory. *MIS Quarterly*, 38(2), 497–520. <https://doi.org/10.25300/MISQ/2014/38.2.08>
- [17] Hofstede, G. (1980). *Culture's consequences: International differences in work-related values*. Sage.
- [18] Hunter, J. E., & Schmidt, F. L. (2004). *Methods of meta-analysis: Correcting error and bias in research findings* (2nd ed.). Sage.

- [19] Keith, M. J., Thompson, S. C., Hale, J., Lowry, P. B., & Greer, C. (2013). Information disclosure on mobile devices: Re-examining privacy calculus with actual user behavior. *International Journal of Human-Computer Studies*, 71(12), 1163–1173. <https://doi.org/10.1016/j.ijhcs.2013.08.016>
- [20] Kietzmann, J., Paschen, J., & Treen, E. (2018). Artificial intelligence in advertising: How marketers can leverage artificial intelligence along the consumer journey. *Journal of Advertising Research*, 58(3), 263–267. <https://doi.org/10.2501/JAR-2018-035>
- [21] Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64, 122–134. <https://doi.org/10.1016/j.cose.2015.07.002>
- [22] Komiak, S. Y. X., & Benbasat, I. (2006). The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS Quarterly*, 30(4), 941–960. <https://doi.org/10.2307/25148760>
- [23] Kumar, V., Rajan, B., Gupta, S., & Dalla Pozza, I. (2019). Customer engagement in service. *Journal of the Academy of Marketing Science*, 47(1), 138–160. <https://doi.org/10.1007/s11747-017-0565-2>
- [24] Laufer, R. S., & Wolfe, M. (1977). Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of Social Issues*, 33(3), 22–42. <https://doi.org/10.1111/j.1540-4560.1977.tb01880.x>
- [25] Li, Y. (2011). Empirical studies on online information privacy concerns: Literature review and an integrative framework. *Communications of the Association for Information Systems*, 28(1), 453–496. <https://doi.org/10.17705/1CAIS.02828>
- [26] Liang, T. P., Lai, H. J., & Ku, Y. C. (2011). Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings. *Journal of Management Information Systems*, 23(3), 45–70. <https://doi.org/10.2753/MIS0742-1222230303>
- [27] Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4), 336–355. <https://doi.org/10.1287/isre.1040.0032>
- [28] Martin, K. D., Borah, A., & Palmatier, R. W. (2017). Data privacy: Effects on customer and firm performance. *Journal of Marketing*, 81(1), 36–58. <https://doi.org/10.1509/jm.15.0497>
- [29] Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709–734. <https://doi.org/10.5465/amr.1995.9508080335>
- [30] McDonald, A. M., & Cranor, L. F. (2008). The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 4(3), 543–568.
- [31] McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). Developing and validating trust measures for e-commerce: An integrative typology. *Information Systems Research*, 13(3), 334–359. <https://doi.org/10.1287/isre.13.3.334.81>
- [32] Milberg, S. J., Burke, S. J., Smith, H. J., & Kallman, E. A. (1995). Values, personal information privacy, and regulatory approaches. *Communications of the ACM*, 38(12), 65–74. <https://doi.org/10.1145/219663.219683>
- [33] Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., & The PRISMA Group. (2009). Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. *PLOS Medicine*, 6(7), e1000097. <https://doi.org/10.1371/journal.pmed.1000097>
- [34] Morgan, R. M., & Hunt, S. D. (1994). The commitment-trust theory of relationship marketing. *Journal of Marketing*, 58(3), 20–38. <https://doi.org/10.1177/002224299405800302>
- [35] Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1), 100–126. <https://doi.org/10.1111/j.1745-6606.2006.00070.x>

- [36] Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- [37] Pavlou, P. A. (2003). Consumer acceptance of electronic commerce: Integrating trust and risk with the technology acceptance model. *International Journal of Electronic Commerce*, 7(3), 101–134. <https://doi.org/10.1080/10864415.2003.11044275>
- [38] Peterson, R. A., & Brown, S. P. (2005). On the use of beta coefficients in meta-analysis. *Journal of Applied Psychology*, 90(1), 175–181. <https://doi.org/10.1037/0021-9010.90.1.175>
- [39] Petronio, S. (2002). *Boundaries of privacy: Dialectics of disclosure*. State University of New York Press.
- [40] Rust, R. T., & Huang, M. H. (2014). The service revolution and the transformation of marketing science. *Marketing Science*, 33(2), 206–221. <https://doi.org/10.1287/mksc.2013.0836>
- [41] Sheng, H., Nah, F. F. H., & Siau, K. (2008). An experimental study on ubiquitous commerce adoption: Impact of personalization and privacy concerns. *Journal of the Association for Information Systems*, 9(6), 344–376. <https://doi.org/10.17705/1jais.00160>
- [42] Smit, E. G., van Noort, G., & Voorveld, H. A. M. (2014). Understanding online behavioural advertising: User knowledge, privacy concerns and online coping behaviour in Europe. *Computers in Human Behavior*, 32, 15–22. <https://doi.org/10.1016/j.chb.2013.11.008>
- [43] Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4), 989–1015. <https://doi.org/10.2307/41409970>
- [44] Smith, H. J., Milberg, S. J., & Burke, S. J. (1996). Information privacy: Measuring individuals' concerns about organizational practices. *MIS Quarterly*, 20(2), 167–196. <https://doi.org/10.2307/249477>
- [45] Tam, K. Y., & Ho, S. Y. (2006). Understanding the impact of web personalization on user information processing and decision outcomes. *MIS Quarterly*, 30(4), 865–890. <https://doi.org/10.2307/25148757>
- [46] Turow, J., King, J., Hoofnagle, C. J., Bleakley, A., & Hennessy, M. (2009). Americans reject tailored advertising and three activities that enable it. SSRN Working Paper. <https://doi.org/10.2139/ssrn.1478214>
- [47] Van Doorn, J., & Hoekstra, J. C. (2013). Customization of online advertising: The role of intrusiveness. *Marketing Letters*, 24(4), 339–351. <https://doi.org/10.1007/s11002-012-9222-1>
- [48] Vesanen, J. (2007). What is personalization? A conceptual framework. *European Journal of Marketing*, 41(5/6), 409–418. <https://doi.org/10.1108/03090560710737534>
- [49] Westin, A. F. (1967). *Privacy and freedom*. Atheneum.
- [50] White, T. B., Zahay, D. L., Thorbjornsen, H., & Shavitt, S. (2008). Getting too personal: Reactance to highly personalized email solicitations. *Marketing Letters*, 19(1), 39–50. <https://doi.org/10.1007/s11002-007-9020-0>
- [51] Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2023). 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>
- [52] Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.