

RESEARCH ARTICLE

The AI Trust Gap: Consumer Perceptions of AI-Powered Digital Marketing

Jennifer Wehr Holt¹ , Linh Nguyen² ¹ Department of Management, Haile College of Business, Northern Kentucky University, USA, ² Department of Marketing, Sports Business and Construction Management, Haile College of Business, Northern Kentucky University, USA

Keywords: artificial intelligence, digital marketing, consumer trust, personalization, privacy, transparency, algorithm aversion, AI marketing

<https://doi.org/10.63963/001c.172170>

Marketing Management Journal

Vol. 36, Issue 1, 2026

Artificial intelligence (AI) is increasingly used in digital marketing to personalize communications, recommend products, and automate customer interactions. Despite these capabilities, consumer trust in AI-powered marketing remains uncertain. This exploratory study examines consumer perceptions of AI-driven digital marketing using survey data collected from 500 U.S. adults. Findings suggest that consumers generally recognize the usefulness of AI-powered marketing but remain cautious about whether organizations using AI act in consumers' best interests. Respondents expressed concerns regarding privacy, data use, and organizational trustworthiness, while transparency and data security emerged as important trust-building factors. Moderate daily AI users (2–5 hours) reported higher levels of trust in companies employing AI-driven marketing practices than light users (<2 hours). Overall, the findings suggest that consumer acceptance of AI's functional benefits does not necessarily translate into trust in organizations that deploy these technologies. Implications for marketing managers and directions for future research are discussed.

1. Introduction

Artificial intelligence (AI) has become a defining feature of contemporary digital marketing. Organizations increasingly rely on AI-driven systems to personalize advertising, recommend products, automate customer service interactions, and predict consumer preferences at scale (Davenport et al., 2020). These capabilities allow marketers to process behavioral data with a level of precision that was previously unattainable.

Despite these advances, the growing presence of AI in marketing has generated significant consumer anxiety regarding privacy, transparency, and ethical data use. Consumers are increasingly aware that AI systems collect and analyze personal information to influence purchasing behavior, often through processes that are unclear and difficult to evaluate. Trust, defined as the expectation that exchange partners will act reliably and benevolently, has emerged as a central variable in understanding consumer responses to AI-powered marketing (Morgan & Hunt, 1994).

In AI-mediated marketing environments, trust is not limited to organizational credibility. Consumers must also evaluate the reliability and fairness of algorithmic systems that operate invisibly behind personalization interfaces (Marthews & Tucker, 2019). This opacity creates a fundamental

tension: personalization systems designed to improve marketing relevance may simultaneously generate perceptions of surveillance and manipulation, thereby undermining the trust required for long-term consumer relationships.

The present study makes a focused empirical contribution by documenting what we term the AI trust gap: the disconnect between consumers who accept AI marketing as functionally useful but remain skeptical of the organizations deploying it. Prior research has examined privacy concerns, algorithm aversion, and trust in automation separately; this study brings them together in a single marketing context to show how these forces combine to produce cautious rather than trusting consumer postures toward AI-powered marketing. The study uses statistical comparisons to identify whether consumer trust attitudes lean positive or negative, whether AI adoption changes perceptions of companies across specific dimensions, and how key attitudes relate to one another. Rather than testing a fully specified causal model, the study aims to characterize the trust landscape and generate hypotheses for future research.

The study addresses the following research questions:

RQ1: Do consumer trust attitudes toward AI-powered marketing lean more positive or negative?

RQ2: Does AI adoption influence perceptions of companies?

RQ3: How are key consumer attitudes toward AI-powered marketing related to one another?

RQ4: What broader patterns can be used to characterize the overall AI trust landscape?

2. Literature Review

The following review draws on four streams of relevant theory and research. Rather than presenting these as parallel frameworks, we integrate them around a central question: why might consumers accept AI marketing as useful while remaining skeptical of the organizations using it? Privacy calculus theory explains the cost-benefit logic governing data sharing. The personalization paradox and creepy marketing research show how data-driven personalization can backfire. Trust-in-automation theory explains how familiarity with technology builds confidence over time. Commitment-Trust Theory grounds the organizational trust dimension. Together, these perspectives set up the AI trust gap this study investigates.

2.1. Consumer Trust and Commitment-Trust Theory

Morgan and Hunt's (1994) Commitment-Trust Theory identifies trust as central to cooperative exchange, arguing that parties who trust one another are more likely to commit to long-term relationships and less likely to engage in opportunistic behavior. In digital environments, where consumers cannot directly observe organizational practices, trust plays an especially critical role in reducing perceived risk and sustaining engagement.

In AI-enabled marketing contexts, trust extends beyond the organization itself to encompass confidence in the technological systems generating recommendations and automating interactions. Consumers must develop trust in AI systems they cannot inspect, making transparency and perceived benevolence especially consequential (Murtarelli et al., 2021). Perceived trustworthiness predicts consumer willingness to engage with personalized services and share personal information (Kim & Park, 2013).

Trust has also been shown to shape important marketplace outcomes beyond attitudes and behavioral intentions. Chaudhuri and Ligas (2016) found that trust, rather than satisfaction, was the mechanism through which favorable perceptions of merchandise value translated into consumers' willingness to pay premium prices. Their findings position trust as a relational pathway connecting positive evaluations of a firm to concrete economic behavior. Extending this logic to AI-powered marketing, trust may similarly determine whether favorable perceptions of AI-driven personalization translate into deeper consumer engagement and sustained relationships.

2.2. Trust in Automation and Artificial Intelligence

Trust in AI can be understood as a specialized form of trust in automated systems more broadly. Hoff and Bashir (2015) propose that trust in automation develops through the interaction of dispositional tendencies, situational context, and accumulated experience with the technology itself. Their framework suggests that people evaluate automated systems not only according to objective performance but also according to perceived reliability, predictability, transparency, and the degree of control retained by the user. Because consumers rarely observe the algorithms responsible for personalization and recommendation directly, trust becomes an especially important mechanism for managing the uncertainty inherent in these interactions.

A broad review of the AI trust literature by Glikson and Woolley (2020) identifies trust as one of the most consistent predictors of AI acceptance across contexts. Their synthesis finds that trust in AI is shaped by perceived fairness, explainability, competence, and accountability, while concerns about bias, privacy intrusion, and insufficient transparency tend to erode it. Together, these findings suggest that consumer trust in AI-powered marketing depends not only on technological performance but also on broader perceptions of how, and for what purpose, AI systems are deployed.

2.3. Privacy Calculus and Data-Sharing Behavior

Privacy calculus theory proposes that individuals weigh potential benefits against privacy risks when deciding whether to disclose personal information (Dinev & Hart, 2006). In AI-driven personalization contexts, this calculus is especially salient: personalized marketing may improve relevance and convenience, but it depends on data collection practices that many consumers experience as intrusive or threatening.

This tension has been described in the marketing literature as the personalization paradox. Aguirre et al. (2015) found that while consumers value the relevance and convenience that personalization provides, they simultaneously grow uncomfortable when firms appear to know an excessive amount about them. Their findings indicate that personalization tends to strengthen consumer experience when data practices are transparent and expected, but undermines trust when consumers perceive the underlying data collection as surveillance rather than service.

Related concerns have been explored through the concept of creepy marketing. Moore et al. (2015) argue that consumers label marketing practices as creepy when firms use personal information in ways that exceed what consumers consider socially normative. Their qualitative work identifies three underlying dimensions of this reaction: a sense of privacy invasion, stalking-like surveillance, and violations of broader social norms. These dimensions are particularly relevant to AI-driven marketing, where sophisticated personalization systems can easily create the impression that a firm knows too much, even when no rule has technically been broken.

The benefit-risk tradeoff is context-sensitive and individual-specific (Awad & Krishnan, 2006). Concerns about data misuse, algorithmic opacity, and lack of consumer control have been shown to reduce willingness to disclose personal information even when material benefits are offered (Martin et al., 2017).

2.4. AI-Specific Trust Challenges

Algorithm aversion, the tendency to reject algorithmic recommendations even when they outperform human judgment, is well documented in the behavioral literature (Dietvorst et al., 2015). In marketing and decision-making contexts, algorithm aversion may manifest as skepticism toward AI-generated recommendations, resistance to personalization systems, or general wariness of AI-powered interactions regardless of their functional quality (Jussupow et al., 2024; Liu et al., 2023; Yalcin et al., 2022).

Transparency has emerged as a key factor shaping AI trust. Puntoni et al. (2021) show that consumers approach AI interactions through an experiential lens, evaluating not only outcomes but the process through which AI arrives at them. Opaque machine-learning systems that cannot explain their recommendations create a transparency gap that undermines confidence even when outcomes are satisfactory. This transparency gap motivates the present investigation.

While neither Banerjee et al. (2025) nor Dang et al. (2025) directly examined AI-powered marketing, their findings offer conceptual extensions relevant to this context. Banerjee et al. (2025) found that perceived authenticity shapes consumer engagement with social media influencer content, and Dang et al. (2025) found that trending content elements in short-form advertisements generated greater engagement than content without such features. By extension — and not as a direct finding of those

studies — these results suggest that consumers may similarly evaluate AI-generated marketing content not only on informational value but also on perceived authenticity and fit with platform-specific communication norms.

3. Methodology

3.1. Research Design and Orientation

This study employed a quantitative, cross-sectional survey design to examine consumer attitudes toward AI-powered digital marketing. Consistent with the objectives of exploratory research, the study seeks to identify directional patterns and relationships that may inform future theory development rather than test a fully specified causal model (Stebbins, 2001). Because cross-sectional survey designs provide limited evidence regarding temporal ordering and causality, the findings should be interpreted as hypothesis-generating rather than confirmatory (Rindfleisch et al., 2008).

3.2. Sample and Data Collection

Data were collected through a nationally distributed online consumer panel administered by Qualtrics. The final analytical sample consisted of 500 U.S. adult consumers who completed the full survey (N = 553 initiated; 53 excluded for incomplete responses). All respondents confirmed familiarity with digital marketing and awareness of AI prior to entering the main survey. An English-language fluency check required participants to compose a written description of a product displayed on screen; respondents who failed were excluded. Demographic composition: gender was approximately equal (48.6% male, 51.4% female); age ranged across all adult cohorts with the largest segment aged 55–64 (23.0%); racial composition was 64.8% Caucasian, 13.4% African American, 8.8% Hispanic, and 6.8% Asian; the majority held at least a college degree.

One limitation of the survey design is that respondents were asked about ‘AI-powered digital marketing’ generally, rather than a specific AI application. Different respondents may have had different AI applications in mind when answering, which introduces heterogeneity in the referent and should be considered when interpreting the findings.

3.3. Measures

The survey instrument was developed based on constructs from the consumer trust and AI acceptance literatures. Items were adapted from previously validated scales where available; study-specific items were developed for constructs without established measures. All Likert-scale items used a 7-point format (1 = Strongly Disagree to 7 = Strongly Agree). [Table 1](#) summarizes the constructs, sample items, item counts, scale sources, and reliability coefficients. A complete listing of all survey items is provided in Appendix A.

Table 1. Survey Constructs, Sources, and Reliability (N = 500)

Construct	Sample Item	# Items	α / SB
Individual Trust (Q2) (Items adapted from Glaeser et al., 2000; McKnight et al., 2002)	I am a trustworthy person; I trust others; I am an ethical person; I trust AI	4 items	.688*
AI DM Attitude (Q18) (Adapted from Bruner & Kumar, 2000)	The use of AI for personalized digital marketing is [good / useful / sensible / responsible]	4 items	.855
General Company Trust (Q19_1) (Adapted from Flavián et al., 2006; Mayer et al., 1995)	Companies openly share information (Transparency); Companies demonstrate competence (Competence) [7 dimensions]	7 items	.870
AI Company Trust (Q19_2) (Same items as Q19_1, reframed for AI-using companies)	Same 7 dimensions, framed for companies that use AI for digital marketing	7 items	.894
AI Benevolence (Q20) (Study-developed item)	I believe when companies use AI to create personalized digital marketing for me, this company has my best interests at heart	1 item	—
AI Trust Impact (Q12/Q13) (Study-developed items)	Do companies using AI (or personalized AI marketing) increase or decrease your trust?	2 items	.933 SB
Data Sharing Willingness (Q15/Q16) (Adapted from Malhotra et al., 2004)	Are you willing to accept risks of giving personal information to receive benefits from companies?	2 items	.953 SB
Trust-Building Preferences (Q17) (Study-developed)	What can companies do to increase your trust in AI marketing? (select all that apply)	5 options	—
Demographics (S1, S2, D2)	Gender; age group; daily hours of AI use	3 items	—

Note. α = Cronbach's alpha; SB = Spearman-Brown coefficient. All Likert scales anchored at 1 = Strongly Disagree to 7 = Strongly Agree. *Individual Trust items are analyzed separately rather than as a composite due to low inter-item consistency, consistent with reviewer guidance. A complete survey instrument is provided in Appendix A.

3.4. Data Preparation

Prior to analysis, a data quality check identified a coding error in the individual trust items (Q2): the data export had recorded a value of 9 in place of the intended value of 7 (Strongly Agree). This was identified by examining the range of values against the documented 1–7 scale and confirmed via the value label mapping in the SPSS export file. All instances were corrected prior to analysis. No other variables showed irregularities.

3.5. Reliability

Internal consistency reliability was assessed using Cronbach's alpha for multi-item scales and Spearman-Brown coefficients for two-item measures. The AI Digital Marketing Attitude scale demonstrated good reliability (α = .855), while the General Company Trust (α = .870) and AI Company Trust (α = .894) scales demonstrated strong reliability. The Individual Trust measure exhibited lower internal consistency (α = .688), consistent with reviewers' observation that the four items (self-trustworthiness, interpersonal trust, personal ethics, and trust in AI) may capture related but conceptually distinct dimensions. Accordingly, these items are analyzed and interpreted individually rather than as a composite. For the two-item measures, AI Trust Impact demonstrated excellent reliability (Spearman-Brown = .933) and Data Sharing Willingness also demonstrated excellent reliability (Spearman-Brown = .953).

Table 2. Average Scores and One-Sample t-Tests vs. Scale Midpoint of 4.0 (N = 500)

Variable	M	SD	N	t	p	vs. Midpoint
Individual Trust Items (analyzed separately due to low inter-item consistency; see Table 1)						
Self-trustworthiness	6.23	0.94	500	53.22	<.001	Above***
Trust in others	4.72	1.45	500	11.07	<.001	Above***
Personal ethics	6.07	1.01	500	45.75	<.001	Above***
Trust in AI	4.24	1.49	500	3.60	.001	Above**
Attitude Toward AI in Digital Marketing (Q18, $\alpha = .855$)						
AI DM: good	4.18	1.47	500	2.74	.006	Above**
AI DM: useful	4.87	1.38	500	14.17	<.001	Above***
AI DM: sensible	4.63	1.38	500	10.25	<.001	Above***
AI DM: responsible	4.06	1.36	500	0.92	.358	ns
Outcome Variables						
AI has my best interests (Q20)	3.26	1.60	500	-10.26	<.001	Below***
AI increases trust (Q12)	3.66	1.19	500	-6.49	<.001	Below***
Personalized AI DM increases trust (Q13)	3.62	1.23	500	-6.82	<.001	Below***
Willing to share data for benefits (Q15)	3.40	1.64	496	-8.12	<.001	Below***
Willing to share more data for more benefits (Q16)	3.33	1.71	496	-8.80	<.001	Below***

Note. One-sample t-tests against scale midpoint of 4.0. * $p < .05$, ** $p < .01$, *** $p < .001$. ns = not significant.

3.6. Analytical Approach

Analyses were conducted in Python using the SciPy and Statsmodels libraries. One-sample t-tests compared each construct's average against the scale midpoint of 4.0 (addressing RQ1). Paired-samples t-tests compared how the same respondents rated general companies versus AI-using companies across seven trust dimensions, with Bonferroni correction applied for multiple comparisons (addressing RQ2). Cohen's d and 95% confidence intervals are reported alongside significance tests to communicate practical effect magnitude. Pearson correlations and a multiple linear regression model, including variance inflation factor (VIF) diagnostics to assess multicollinearity, examined attitudinal interrelationships and predictors of AI company trust (addressing RQ3). One-way ANOVAs with Tukey HSD post hoc tests examined whether trust differed across subgroups defined by AI usage frequency, age, and gender, alongside descriptive privacy analyses (addressing RQ4). Common method bias was assessed via Harman's single-factor test applied to the set of primary perceptual variables.

4. Results

4.1. Research Question 1: Do Consumer Trust Attitudes Toward AI-Powered Marketing Lean More Positive or Negative?

To address RQ1, one-sample t-tests compared respondents' average ratings against the scale midpoint of 4.0. [Table 2](#) and [Figure 1](#) present the results.




Figure 1. Individual Trust Profile, Perceptions of Companies, and Attitudes Toward AI in Digital Marketing

Note. Bar heights represent average scores (7-point scale). Error bars = SD. Dashed line = scale midpoint (4.0). *** $p < .001$, ** $p < .01$.

Consumers reported high self-trustworthiness ($M = 6.23$, $SD = 0.94$) and personal ethics ($M = 6.07$, $SD = 1.01$), both significantly above the midpoint ($p < .001$). Trust in others was also above neutral ($M = 4.72$, $SD = 1.45$, $p < .001$). Trust in AI was above the midpoint ($M = 4.24$, $SD = 1.49$, $p = .001$), though only modestly so, suggesting cautious rather than confident trust.

Attitudes toward AI in digital marketing were generally favorable. Respondents rated AI DM as useful ($M = 4.87$, $p < .001$) and sensible ($M = 4.63$, $p < .001$), while perceptions that AI DM is responsible did not differ significantly from neutral ($M = 4.06$, $p = .358$). This pattern suggests consumers recognize practical benefits of AI marketing while remaining uncertain about its ethical deployment.

In contrast, several outcome measures were significantly below neutral. Respondents disagreed that AI-using companies act in their best interests ($M = 3.26$, $SD = 1.60$, $p < .001$), and reported that both AI use generally ($M = 3.66$) and personalized AI marketing specifically ($M = 3.62$) decrease rather than increase trust. Willingness to share data for benefits was also below neutral ($M = 3.40$ and $M = 3.33$). Taken together, these findings suggest that consumer trust toward AI-powered marketing is best characterized as cautious and conditional. Consumers accept AI's functional utility while remaining skeptical about the organizations deploying it — precisely the AI trust gap this study investigates.

4.2. Research Question 2: Does AI Adoption Influence Perceptions of Companies?

To address RQ2, paired-samples t -tests compared how the same 500 respondents rated general companies versus AI-using companies across seven trust dimensions. Bonferroni correction was applied given the seven simultaneous comparisons (corrected $\alpha = .007$). [Table 3](#) and [Figure 2](#) present the results, including effect sizes (Cohen's d) and 95% confidence intervals.

Table 3. Paired-Samples t-Tests: General vs. AI-Using Company Trust (N = 500)

Dimension	General M (SD)	AI M (SD)	t	p	Sig.	d	95% CI
Humanity	3.57 (1.57)	3.56 (1.54)	0.14	.887	ns	0.006	[-0.077, 0.089]
Transparency	3.48 (1.56)	3.56 (1.57)	-1.59	.113	ns	-0.071	[-0.170, 0.018]
Capability	4.72 (1.22)	4.56 (1.40)	3.24	.001	**	0.145	[0.062, 0.250]
Credibility	3.95 (1.40)	3.86 (1.51)	1.97	.049	ns†	0.088	[0.001, 0.183]
Reliability	4.19 (1.35)	4.12 (1.43)	1.60	.110	ns	0.072	[-0.016, 0.160]
Integrity	3.64 (1.49)	3.58 (1.55)	1.52	.129	ns	0.068	[-0.018, 0.146]
Competence	5.16 (1.09)	4.92 (1.30)	5.03	<.001	***	0.225	[0.144, 0.328]

Note. Paired-samples t-tests (N = 500, same respondents rated both conditions). d = Cohen's d for paired design (mean difference / SD of differences). 95% CI = confidence interval around the mean difference. Bonferroni-corrected $\alpha = .007$ (seven comparisons); † Credibility (p = .049) is not significant after correction. Effect sizes: d < 0.20 = negligible, d = 0.20–0.49 = small, d = 0.50–0.79 = medium. *p < .05, **p < .01, ***p < .001.




Figure 2. Comparing Consumer Trust Ratings: General vs. AI-Using Companies

Note. Paired-samples t-tests (N = 500). Error bars = SD. Dashed line = scale midpoint (4.0). ** p < .01, *** p < .001 (Bonferroni-corrected).

After Bonferroni correction, two dimensions showed statistically significant differences. Competence was rated meaningfully lower for AI-using companies (M = 4.92) than for companies generally (M = 5.16), $t(499) = 5.03$, $p < .001$, $d = 0.225$, 95% CI [0.144, 0.328]. Capability also differed significantly (General M = 4.72; AI M = 4.56), $t(499) = 3.24$, $p = .001$, $d = 0.145$, 95% CI [0.062, 0.250]. Credibility (p = .049) did not survive Bonferroni correction and should be interpreted with caution.

No significant differences after correction emerged on human-centered dimensions: Humanity (p = .887), Transparency (p = .113), Integrity (p = .129), and Reliability (p = .110). These effect sizes are small in magnitude ($d < 0.20$ for all non-significant dimensions), and it is important to note that non-significant results do not establish equivalence between the two conditions. Rather, the data do not provide sufficient evidence to conclude these perceptions differ.

Table 4. Pearson Correlations Among Key Trust and Attitude Variables (N = 496–500)

Variable Pair	r	p	Relationship
Trust in AI × AI DM Attitude	0.649	<.001	Moderate-strong, positive
AI DM Attitude × AI Company Trust	0.720	<.001	Strong, positive
General Company Trust × AI Company Trust	0.874	<.001	Very strong, positive
AI Best Interests (Q20) × AI Company Trust	0.709	<.001	Strong, positive
Data Sharing Willingness (Q15) × AI Company Trust	0.410	<.001	Moderate, positive

Note. N = 496–500. Pearson r. All correlations significant at $p < .001$. The strong correlation between General Company Trust and AI Company Trust ($r = .874$) was examined for common method bias via Harman's single-factor test; the first unrotated factor explained 65.8% of total variance, warranting cautious interpretation of the magnitude of shared relationships. See Limitations.

Table 5. Multiple Regression Predicting AI Company Trust (N = 496)

Predictor	B	SE	t	p	Sig.	VIF
(Constant)	0.884	0.125	7.08	<.001	***	—
Trust in AI	0.053	0.030	1.75	.081	ns	1.92
AI DM Attitude	0.334	0.043	7.78	<.001	***	2.76
AI Has Best Interests (Q20)	0.283	0.028	10.25	<.001	***	1.86
AI Increases Trust (Q12)	0.134	0.043	3.14	.002	**	2.41
Willing to Share Data (Q15)	0.007	0.023	0.29	.774	ns	1.38

Note. Dependent variable: AI Company Trust composite. N = 496. $R^2 = .645$, Adjusted $R^2 = .641$, $F(5,490) = 177.76$, $p < .001$. B = unstandardized coefficient; SE = standard error; VIF = variance inflation factor. All VIFs < 3.0, indicating acceptable multicollinearity. ** $p < .01$, *** $p < .001$.

The observed pattern suggests a selective competence-related penalty associated with AI adoption: companies using AI are perceived as less capable and less competent than companies generally, but there is no significant evidence of differences on ethical or relational dimensions. RQ2 is answered: AI adoption is associated with selectively lower perceptions of competence-related trust dimensions, rather than a broad or uniform erosion of trust.

4.3. Research Question 3: How Are Key Consumer Attitudes Related to One Another?

To address RQ3, Pearson correlations characterized relationships among key constructs, and a multiple regression identified the strongest attitudinal predictors of AI company trust. [Table 4](#) presents correlations and [Table 5](#) presents regression results.

All relationships were positive and significant at $p < .001$. The strongest was between General Company Trust and AI Company Trust ($r = .874$), indicating that pre-existing corporate trust attitudes carry strongly into AI-specific contexts. Attitude toward AI DM was also strongly related to AI company trust ($r = .720$). Harman's single-factor test indicated that the first unrotated factor explained 65.8% of total variance, which exceeds the 50% threshold used as a heuristic for potential common method bias. This result does not necessarily invalidate the findings but warrants caution when interpreting the magnitude of shared relationships, and is acknowledged as a limitation of the single-source survey design.

Table 6. Group Comparisons of AI Company Trust by Usage, Age, and Gender

Grouping Variable	Groups	F / t	df	p	Sig.
Daily AI usage (D2)	<2 hrs (n=407) / 2–5 hrs (n=70) / 6–10 hrs (n=10) / >10 hrs (n=6)	6.15	3, 489	<.001	***
Age group (S2)	18–24 through Over 65 (6 groups)	1.48	5, 494	.196	ns
Gender (S1)	Male vs. Female	0.96 ¹	498	.340	ns

Note. Dependent variable: AI Company Trust composite. One-way ANOVAs for usage and age; independent-samples t-test for gender (¹t-value). Tukey HSD post hoc tests revealed that only the <2 hrs and 2–5 hrs groups differed significantly ($p = .001$); no significant differences were found involving the 6–10 hrs ($n = 10$) or >10 hrs ($n = 6$) groups. Caution is warranted regarding the two small groups given insufficient statistical power. *** $p < .001$.

The overall model was significant and explained substantial variance, $R^2 = .645$, $F(5, 490) = 177.76$, $p < .001$. VIF scores ranged from 1.38 to 2.76, all well below the conventional concern threshold of 10 (or conservative threshold of 5), indicating acceptable multicollinearity. Attitude toward AI DM ($B = .334$, $p < .001$) and belief that AI acts in consumers' best interests ($B = .283$, $p < .001$) were the strongest predictors. Perceived trust impact of AI also contributed significantly ($B = .134$, $p = .002$). General trust in AI ($p = .081$) and data-sharing willingness ($p = .774$) were not significant once other variables were controlled. It is important to note that the predictor 'AI increases trust' and the outcome 'AI company trust' are conceptually related, and this overlap may contribute to the predictive strength of the model. The findings should be interpreted as associative rather than causal, and the regression is best understood as characterizing which attitudes are most closely interrelated.

4.4. Research Question 4: What Broader Patterns Characterize the AI Trust Landscape?

To address RQ4, one-way ANOVAs with Tukey HSD post hoc tests examined whether AI company trust differed across consumer subgroups, and descriptive analyses characterized privacy concerns and trust-building preferences. [Table 6](#) and [Figure 3](#) present the group-comparison results.

The overall ANOVA for daily AI usage was significant, $F(3, 489) = 6.15$, $p < .001$. Tukey HSD post hoc tests revealed that only the difference between the less-than-2-hours group ($M = 3.93$, $n = 407$) and the 2-to-5-hours group ($M = 4.51$, $n = 70$) was statistically significant ($p = .001$). No other group comparisons reached significance, including those involving the 6-to-10-hours ($n = 10$) and greater-than-10-hours ($n = 6$) groups. Given the extremely small sample sizes in these latter two groups, conclusions about their means should not be drawn, and future research should examine high-usage consumers with an appropriately powered design.

Age group, $F(5, 494) = 1.48$, $p = .196$, and gender, $t(498) = 0.96$, $p = .340$, were not significantly associated with AI company trust. These results should not be interpreted as establishing that age and gender are unrelated to trust; rather, this study did not find sufficient evidence of a relationship in this sample.

Tables [7](#) and [8](#) present privacy and trust-building preference findings.




Figure 3. AI Company Trust by Daily AI Usage Frequency

Note. One-way ANOVA, $F(3, 489) = 6.15, p < .001$. Error bars = SD. Tukey HSD: only <2 hrs vs. 2–5 hrs significant ($p = .001$).

Table 7. Consumer Trust Response to Company Knowledge of Personal Data (N = 500)

Response	n	%
Greater company data knowledge increases trust	73	14.6%
Greater company data knowledge decreases trust	220	44.0%
Greater company data knowledge has no impact on trust	207	41.4%

Note. N = 500.

Table 8. Consumer Trust-Building Action Preferences (N = 499, Multiple Response)

Trust-Building Action	n	% of Respondents
Ensure data security (Q17_3)	436	87.4%
Be transparent about data usage (Q17_1)	414	83.0%
Offer opt-out options (Q17_4)	390	78.2%
Provide clear privacy policies (Q17_2)	378	75.8%
Provide more value in exchange for data (Q17_5)	239	47.9%

Note. N = 499. Multiple responses permitted; percentages reflect the proportion of respondents selecting each option and may be compared across options but do not sum to 100%.

A plurality of respondents (44.0%, $n = 220$) reported that greater company knowledge of their personal data decreases trust, while only 14.6% reported increased trust. When asked what would build trust in AI-driven marketing, consumers most frequently selected ensuring data security (87.4%), followed by transparency about data usage (83.0%), opt-out options (78.2%), clear privacy policies (75.8%), and providing more value in exchange for data




Figure 4. Privacy Perceptions and Trust-Building Preferences (Corrected)

Note. Panel A: N = 500. Panel B: N = 499, multiple responses permitted. Percentages reflect proportion of respondents selecting each option.

(47.9%). Note that this last figure (47.9%) differs from the value reported in a prior version of the manuscript (64.7%), which was the result of a data extraction error; Table 8 reflects the verified counts from the final cleaned dataset.

5. Discussion

This study examined consumer trust in AI-powered digital marketing through an exploratory survey of 500 U.S. adults. The findings are organized around each research question and connected to the theoretical framework introduced in the literature review.

5.1. Research Question 1

The findings associated with RQ1 reveal the core AI trust gap: consumers accept AI marketing as functionally useful while remaining skeptical of the organizations deploying it. Trust in AI was only marginally above neutral, and outcome measures — particularly whether AI companies act in consumers' best interests — were significantly below neutral. This pattern supports Hoff and Bashir's (2015) and Glikson and Woolley's (2020) observation that functional competence alone does not generate trust when users lack visibility into how systems operate. It also aligns with privacy calculus theory: even when consumers recognize the benefits of personalization, concerns about organizational intentions and data practices suppress trust and data-sharing willingness.

5.2. Research Question 2

The findings associated with RQ2 show that AI adoption is associated with a targeted rather than uniform reduction in company trust. After Bonferroni correction, only Competence and Capability differed significantly between general and AI-using companies, with small effect sizes ($d = 0.145\text{--}0.225$). No significant differences were found on Humanity, Integrity, Transparency, or Reliability. This selective pattern suggests that consumers question the organizational judgment and execution quality behind AI adoption — not the ethics or relational character of the firm. This finding is consistent with Glikson and Woolley's (2020) observation that perceived competence is a separable and critical component of AI trust. Importantly, the non-significant results on human-centered dimensions should not be interpreted as establishing equivalence; they simply indicate insufficient evidence of differences in this sample.

5.3. Research Question 3

The findings associated with RQ3 show that attitudes toward AI-powered marketing are strongly interrelated, and that belief in AI's benevolent intent and overall marketing attitude are most strongly associated with AI company trust. This is consistent with Chaudhuri and Ligas's (2016) finding that trust, rather than general satisfaction, mediates the relationship between positive evaluations and behavioral engagement. Harman's single-factor test result (65.8% of variance explained by the first factor) suggests some degree of common method variance, and readers should be cautious about the magnitude of reported correlations. The regression findings are best understood as identifying which attitudes travel together most closely, rather than establishing causal pathways.

5.4. Research Question 4

The findings associated with RQ4 show that AI usage frequency is statistically associated with AI company trust, but this relationship is more limited than initially suggested. Tukey HSD post hoc tests indicate that only moderate daily AI users (2–5 hours) differ significantly from light users (<2 hours); higher-usage groups are too small to draw conclusions from. This pattern is broadly consistent with Hoff and Bashir's (2015) emphasis on accumulated experience as a key mechanism in the development of automation trust. The privacy findings reinforce the personalization paradox (Aguirre et al., 2015) and creepy marketing concerns (Moore et al., 2015): more company data knowledge tends to decrease rather than increase trust for the plurality of consumers, and value-exchange incentives rank last among trust-building preferences.

6. Managerial Implications

The findings identify several actions marketing managers can take to address the gap between consumers' acceptance of AI's functional benefits and their trust in the organizations deploying it. The results suggest that effective AI implementation requires more than demonstrating technological capability or delivering increasingly personalized experiences. Managers must also consider how AI use affects perceptions of organizational intent, data stewardship, consumer control, and competence.

Demonstrate consumer-oriented intent, not just personalization value. Belief that companies using AI act in consumers' best interests was strongly associated with AI company trust. Managers should therefore communicate not only that AI is being used, but why it is being used and how consumers benefit from it. For example, firms can explain how AI helps reduce irrelevant advertising, improve recommendations, simplify service interactions, or otherwise create identifiable benefits for consumers. This shifts the emphasis from demonstrating technological sophistication to demonstrating that AI serves a legitimate consumer purpose.

Make data security and transparency visible. Consumers identified data security (87.4%) and transparency about data use (83.0%) as the two most frequently selected actions that would increase their trust in AI-driven marketing. Opt-out options (78.2%) and clear privacy policies (75.8%) were also selected substantially more often than providing greater value in exchange for personal data (47.9%). This pattern suggests that trust may be built more effectively through risk reduction, transparency, and consumer control than through attempts to compensate consumers for increased data collection. Managers should therefore make privacy protections visible at the points where data are collected and AI-driven personalization occurs, clearly communicating what information is collected, how it is used, and how it is protected.

Preserve meaningful consumer control. The strong preference for opt-out options suggests that consumers value the ability to control their participation in AI-driven personalization. Managers should make personalization and data-use controls easy to locate, understand, and modify rather than treating them solely as legal or privacy-policy disclosures. Providing consumers with meaningful choices regarding personalization, data use, and AI-mediated interactions may also help reduce perceptions that personalization represents surveillance rather than service. This approach is consistent with the broader privacy findings, in which 44.0% of respondents indicated that greater company knowledge of their personal data decreases trust, compared with only 14.6% who indicated that it increases trust.

Do not assume that AI adoption signals greater organizational competence. AI-using companies were rated significantly lower on competence and capability than companies generally, although the observed effect sizes were small. No significant differences after correction were found on the remaining trust dimensions. For managers, these findings suggest that

adopting AI should not be assumed to enhance consumer perceptions of organizational expertise. Firms may benefit from demonstrating the quality and reliability of AI-supported decisions, maintaining appropriate human oversight, and providing accessible human support when automated systems fail or produce unsatisfactory results. AI implementation should therefore be evaluated not only in terms of efficiency and technological performance but also in terms of how it affects perceptions of the organization using the technology.

Build familiarity among consumers with limited AI experience. Moderate daily AI users (2-5 hours) reported significantly higher AI company trust than light users (<2 hours), although the cross-sectional design does not establish whether greater AI use produces greater trust or whether consumers who are more trusting of AI are more likely to use it. Managers could nevertheless consider introducing consumers to AI through low-risk, easily understood applications before asking them to rely on AI in interactions involving more sensitive information or consequential decisions. Product discovery, basic recommendations, and optional AI-assisted customer service may provide opportunities for consumers to become more familiar with AI-enabled interactions. Because the present findings are associative rather than causal, the effectiveness of such strategies should be tested rather than assumed.

Consider existing brand trust when implementing AI personalization. General company trust and AI company trust were very strongly correlated ($r = .874$), suggesting that consumers' existing perceptions of companies are closely related to how they evaluate those companies in an AI context. However, the magnitude of this relationship should be interpreted cautiously given the evidence of potential common method bias in the study. Managers should therefore view AI trust as part of broader brand and relationship management rather than as an isolated technology issue. Organizations with weaker baseline consumer trust should be particularly cautious about assuming that sophisticated AI personalization will independently overcome existing trust concerns.

Manage the AI trust gap as a strategic marketing issue. Collectively, these findings suggest that the central managerial challenge is not convincing consumers that AI-powered marketing is useful; respondents already recognized its functional benefits. Rather, the challenge is converting that perceived usefulness into trust in the organizations deploying the technology. Managers should therefore evaluate AI marketing initiatives not only through traditional performance measures such as engagement, conversion, personalization effectiveness, and efficiency, but also through their effects on consumer trust. Measures of perceived transparency, data security, consumer control, benevolent intent, and organizational competence can help firms identify whether AI implementation is strengthening or weakening customer relationships. Ultimately, effective AI marketing may depend less on

demonstrating what the technology can do and more on demonstrating that it is being used responsibly, transparently, and in ways that consumers perceive as serving their interests.

7. Limitations and Future Research

This study has several important limitations. First, the cross-sectional design does not support causal inference. The associations between AI usage and trust, and between attitudinal predictors and AI company trust, could reflect selection effects rather than causal relationships. Second, Harman's single-factor test indicated potential common method bias (first factor = 65.8%), a consequence of collecting all perceptual data from the same respondents in the same survey. Future research should consider multi-source designs or procedural remedies such as temporal separation of predictors and outcomes.

Third, the sample is limited to U.S. adults, restricting cross-cultural generalizability. Fourth, the survey asked about 'AI-powered digital marketing' generally, without specifying a particular application; different respondents may have had different technologies in mind, introducing heterogeneity in the referent construct. Fifth, the regression model included only attitudinal predictors; personality traits, prior data breach experiences, and industry-specific factors were not measured.

Sixth, several trust-building preference response categories (Q17) had a small number of additional responses beyond the five main categories (Q17_6: $n = 17$, 3.4%) that were not analyzed in the main body of the paper. Future research should examine the full range of consumer-identified trust-building actions more comprehensively.

Future research could address these limitations through longitudinal or experimental designs that establish causal direction, cross-national comparisons, and studies that examine specific AI applications (e.g., recommendation systems, chatbots, targeted advertising) rather than AI-powered marketing in general. Structural equation modeling could also test the theoretical relationships implied by the present correlation and regression findings.

8. Conclusion

This study documents a clear AI trust gap in consumer perceptions of AI-powered digital marketing: consumers accept the functional usefulness of AI while remaining skeptical about the organizations deploying it. AI-using companies are specifically penalized on competence-related trust dimensions, not on ethical or relational ones. Belief in AI's benevolent intent and overall attitude toward AI marketing are the attitudes most strongly associated with AI company trust, while age and gender show no significant association. Moderate daily AI users report higher trust than light users, suggesting a role for experience in building comfort with AI marketing over time.

These findings suggest that companies hoping to build consumer trust in AI-powered marketing should invest in visible governance, clear communication of benevolent intent, and data security assurances rather than expanded personalization or incentive programs. The gap between AI marketing practice and consumer trust is real, measurable, and — these findings suggest — unlikely to close without deliberate attention to how, not just how well, AI is deployed.

Submitted: August 05, 2026 CDT. Accepted: October 05, 2026 CDT. Published: October 16, 2026 CDT.



This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CCBY-4.0). View this license's legal deed at <http://creativecommons.org/licenses/by/4.0> and legal code at <http://creativecommons.org/licenses/by/4.0/legalcode> for more information.

References

- 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>
- 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>
- Banerjee, N., Rawal, M., Saavedra Torres, J. L., & Bagherzadeh, R. (2025). Beyond the facade: Exploring the authenticity of social media influencers and its influence on consumer brand engagement and advocacy. *Journal of Promotion Management*, 31(7), 1114–1138. <https://doi.org/10.1080/10496491.2025.2571948>
- Bruner, G. C., & Kumar, A. (2000). Web commercials and advertising hierarchy-of-effects. *Journal of Advertising Research*, 40(1–2), 35–42. <https://doi.org/10.2501/jar-40-1-2-35-42>
- Chaudhuri, A., & Ligas, M. (2016). The role of store trust and satisfaction in creating premium prices. *Marketing Management Journal*, 26(1), 1–17. <https://doi.org/10.63963/001c.151124>
- Dang, A., Saravade, S., Spradlin, J., Nworah, P., Linstead, A., & Munger, R. (2025). Should brands utilize trending content elements for their short-form advertisements? *Marketing Management Journal*, 35(1), 1–22. <https://doi.org/10.63963/001c.144766>
- 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>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
- 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>
- Flavián, C., Guinalfú, M., & Gurrea, R. (2006). The role played by perceived usability, satisfaction and consumer trust on website loyalty. *Information & Management*, 43(1), 1–14. <https://doi.org/10.1016/j.im.2005.01.002>
- Glaeser, E. L., Laibson, D. I., Scheinkman, J. A., & Soutter, C. L. (2000). Measuring trust. *Quarterly Journal of Economics*, 115(3), 811–846. <https://doi.org/10.1162/003355300554926>
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. <https://doi.org/10.5465/annals.2018.0057>
- Hoff, K. A., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, 57(3), 407–434. <https://doi.org/10.1177/0018720814547570>
- Jussupow, E., Benbasat, I., & Heinzl, A. (2024). An integrative perspective on algorithm aversion and appreciation in decision-making. *MIS Quarterly*, 48(4), 1575–1590. <https://doi.org/10.25300/MISQ/2024/18512>
- Kim, D. J., & Park, C. H. (2013). The role of trust in e-commerce relational exchange: A unified model. *Information Systems Research*, 20(1), 45–62.

- Liu, N. T. Y., Kirshner, S. N., & Lim, E. T. (2023). Is algorithm aversion WEIRD? A cross-country comparison of individual-differences and algorithm aversion. *Journal of Retailing and Consumer Services*, 72, 103259. <https://doi.org/10.1016/j.jretconser.2023.103259>
- 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>
- Marthews, A., & Tucker, C. E. (2019). Privacy regulation and online advertising. *Management Science*, 65(12), 5669–5688. <https://doi.org/10.1287/mnsc.2018.3273>
- 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>
- 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>
- 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>
- Moore, R. S., Moore, M. L., Shanahan, K. J., Horky, A., & Mack, B. (2015). Creepy marketing: Three dimensions of perceived excessive online privacy violation. *Marketing Management Journal*, 25(1), 42–53. <https://doi.org/10.63963/001c.151082>
- 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>
- Murtarelli, G., Gregory, A., & Romenti, S. (2021). A conversation-based perspective for shaping ethical human–machine interactions. *Journal of Business Research*, 129, 927–935. <https://doi.org/10.1016/j.jbusres.2020.09.018>
- Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing*, 85(1), 131–151. <https://doi.org/10.1177/0022242920953847>
- Rindfleisch, A., Malter, A. J., Ganesan, S., & Moorman, C. (2008). Cross-sectional versus longitudinal survey research: Concepts, findings, and guidelines. *Journal of Marketing Research*, 45(3), 261–279. <https://doi.org/10.1509/jmkr.45.3.261>
- Stebbins, R. A. (2001). *Exploratory research in the social sciences*. SAGE Publications. <https://doi.org/10.4135/9781412984249>
- Yalcin, G., Lim, S., Puntoni, S., & van Osselaer, S. M. (2022). Thumbs up or down: Consumer reactions to decisions by algorithms versus humans. *Journal of Marketing Research*, 59(4), 696–717. <https://doi.org/10.1177/00222437211070016>

Appendix A. Complete Survey Instrument

The following items constitute the complete survey instrument used in this study. All Likert-scale items were measured on a 7-point scale (1 = Strongly Disagree, 7 = Strongly Agree) unless otherwise noted.

Individual Trust (Q2)

- Q2_1: I am a trustworthy person.
- Q2_2: I trust other people.
- Q2_3: I am an ethical person.
- Q2_4: I trust AI (artificial intelligence).

Attitude Toward AI in Digital Marketing (Q18)

The use of AI for personalized digital marketing is...

- Q18_1: Good.
- Q18_2: Useful.
- Q18_3: Sensible.
- Q18_4: Responsible.

General Company Trust (Q19_1) — 7 Dimensions

Q19_1_Humanity: Companies demonstrate concern for the well-being of consumers.

Q19_1_Transparency: Companies openly share information about their practices.

Q19_1_Capability: Companies have the skills and knowledge to meet consumer needs.

Q19_1_Credibility: Companies can be believed when they make claims.

Q19_1_Reliability: Companies consistently deliver on their promises.

Q19_1_Integrity: Companies adhere to ethical principles.

Q19_1_Competence: Companies demonstrate high levels of professional competence.

AI Company Trust (Q19_2) — Same 7 Dimensions, Reframed

Instructions: Please rate the following statements for companies that use AI to create personalized digital marketing.

(Same seven items as Q19_1, with 'Companies' replaced by 'Companies that use AI for digital marketing')

AI Benevolence (Q20)

Q20: I believe when companies use AI to create personalized digital marketing for me, this company has my best interests at heart.

AI Trust Impact (Q12/Q13)

Q12: Using AI (in general) to create marketing for me increases / decreases my trust in a company. (7-point scale: 1 = Strongly Decreases, 7 = Strongly Increases)

Q13: Using AI to create personalized digital marketing for me increases / decreases my trust in a company. (7-point scale: 1 = Strongly Decreases, 7 = Strongly Increases)

Data Sharing Willingness (Q15/Q16)

Q15: I am willing to accept the risks of giving personal information to companies in exchange for benefits. (7-point Likert)

Q16: I am willing to share more personal information with companies if they offer more personalized benefits in return. (7-point Likert)

Trust-Building Preferences (Q17)

Q17: When companies use AI for digital marketing, which of the following would increase your trust? (Select all that apply)

Q17_1: Be transparent about how my data is used.

Q17_2: Provide clear privacy policies.

Q17_3: Ensure the security of my personal data.

Q17_4: Offer opt-out options for AI-driven personalization.

Q17_5: Provide greater value or benefits in exchange for my data.

Q17_6: Other (please specify).

Demographics

S1: Gender (Male / Female / Non-binary / Prefer not to say)

S2: Age group (18–24 / 25–34 / 35–44 / 45–54 / 55–64 / 65 or older)

D2: On average, how many hours per day do you use AI tools? (<2 hrs / 2–5 hrs / 6–10 hrs / >10 hrs)