

Consumer Trust in AI-Generated Advertisements. A Comparative Study of AI and Human-Created Advertising Content

Musawer Hakimi¹, Nazar Mohammad Parsa², Mohammad Nawab Turan³

^a Computer Science Department, Samangan University, Samangan, Afghanistan

^b Computer Science Department, Education Faculty, Faryab University, Faryab, Afghanistan.

^c Computer Science Faculty, Muğla Sıtkı Koçman University, Turkey

Corresponding Author Email Address: musawer@adc.edu.in

ARTICLE INFO

DOI:

<https://doi.org/10.38043/commusty.v5i1.7734>

Article history:

Received: 2026-01-11

Revised: 2026-03-22

Accepted: 15-04-30

Keywords.

AI advertising;

consumer trust;

perceived authenticity;

AI literacy;

structural equation modelling;

technology acceptance

ABSTRACT

The rise of generative artificial intelligence in advertising has created pressing questions about whether consumers can or will trust content they know a machine has crafted. This study examined how 412 adult consumers evaluated AI-generated versus human-created advertisements across six product sectors, investigating the mechanisms of perceived authenticity, cognitive engagement, and emotional response as mediators, and AI literacy as a boundary condition. Using a mixed experimental-survey design followed by structural equation modelling, mediation analysis, and cluster analysis, results showed that AI-generated advertisements received significantly lower trust ratings ($d = 0.89$), yet this gap narrowed substantially among participants with higher AI literacy. Perceived authenticity emerged as the strongest mediator, accounting for 31.2% of the indirect effect. Importantly, trust in AI advertising increased with repeated exposure, suggesting that familiarity attenuates initial scepticism. These findings yield actionable implications for practitioners and advance theoretical understanding of technology-mediated persuasion.

This is an open access article under the [CC-BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.



1. Introduction

The Advertising has always reflected the tools available to those who make it. From hand-painted billboards to television spots filmed over weeks and reviewed by committees, each technological shift has reshaped not only what ads look like but how audiences feel about them. Generative artificial intelligence represents, perhaps, the most disorienting shift yet not because the ads it produces are obviously bad, but because they can be remarkably good. A viewer glancing at an AI-generated image of a smiling family holding a cereal box may notice nothing unusual. Yet if that viewer subsequently learns that the image was machine-made, something changes. The cereal still tastes the same. The family has not become less appealing. And yet the ad, in retrospect, feels different (Dahlén & Edenius, 2007).

What exactly changes and why is the central question of this paper. Consumer trust is a relational construct. It has always depended not only on the message itself but on beliefs about who sent it and with what intention. When a human copywriter crafts an advertisement, there is, at least in theory, a person behind the words who made choices, held opinions, and cared or did not care about the audience. Generative AI complicates this picture in ways that existing advertising theory was not designed to address. The communication lacks a clear author in the conventional sense. The apparent sincerity of the language may be a statistical artefact. The creativity, however striking, emerged from pattern-matching across vast corpora rather than from felt experience (Napoli et al., 2014).

This study begins from the premise that these complications are felt by consumers, even if they struggle to articulate them precisely. We argue that the trust gap between AI-generated and human-created advertising is real, measurable, and theoretically meaningful but that it is neither fixed nor universal. Trust responses depend on psychological mediators, particularly perceptions of authenticity and emotional resonance, and they are shaped by individual-level moderators, especially the degree to which a consumer understands how AI systems work.

The empirical literature on AI-generated content and consumer response is growing rapidly but remains fragmented. Studies have examined individual outcomes click-through rates, brand evaluations, purchase intentions without fully accounting for the psychological processes that connect content source to consumer behaviour (Davenport et al., 2020; Huang & Rust, 2021; Kim & Cheong, 2023). Theoretical frameworks have been borrowed largely from existing advertising and technology acceptance literatures, without sufficient consideration of what is genuinely new about generative AI as a communicative actor (Mick & Fournier, 1998; Sundar & Nass, 2000; Fogg, 2003). Methodologically, most existing work relies on single-exposure experimental designs that cannot capture the longitudinal dynamics of trust formation (Luhmann, 1979; Rousseau et al., 1998).

We address these gaps with a pre-registered mixed-methods study involving 412 participants recruited through Prolific Academic. Participants were exposed to matched pairs of AI-generated and human-created advertisements across six industry sectors, before completing validated measures of trust, authenticity, cognitive engagement, emotional response, purchase intention, brand attitude, and AI literacy. Data were analysed using confirmatory factor analysis, structural equation modelling, bootstrapped mediation analysis, hierarchical regression, and k-means cluster analysis.

The study makes four distinct contributions that advance the existing education–job mismatch and AI advertising literatures in meaningful ways. First, it provides robust evidence for the existence and magnitude of the AI advertising trust deficit using a multi-sector comparative design, extending prior single-sector studies (Kim & Cheong, 2023) and offering the first multi-industry benchmark of the trust gap. Second, it identifies the mediating mechanisms perceived authenticity, cognitive engagement, and emotional response that link content source to trust outcomes, offering a theoretically grounded account of why the gap exists that prior descriptive studies lacked. Third, it demonstrates that AI literacy moderates the trust gap in an asymmetric and counterintuitive fashion: high-literacy consumers trust AI ads more, not less, suggesting that education rather than disclosure alone may be the key practical lever a finding that challenges the prevailing assumption in the literature that greater knowledge of AI increases scepticism. Fourth, it introduces longitudinal exposure data showing that trust in AI advertising increases with repeated exposure, pointing toward familiarity as a trust-building mechanism that practitioners can leverage, a dynamic absent from the predominantly cross-sectional literature to date.

The remainder of this paper is organized as follows. Section 2 reviews the theoretical foundations and develops the research hypotheses. Section 3 describes the research methodology, including sample, stimuli, measures, and analytical procedures. Section 4 presents results. Section 5 discusses findings in relation to theory and practice, and Section 6 concludes with limitations and directions for future research.

2. Theoretical Framework and Hypotheses Development

2.1 Conceptual Framework Overview

The conceptual framework guiding this study integrates four established theoretical traditions to account for the complexity of consumer responses to AI-generated advertising. The Elaboration Likelihood Model (Petty & Cacioppo, 1986) provides the foundational account of how consumers process persuasive messages under varying conditions of motivation and ability. Signalling Theory (Spence, 1973) explains how the source characteristics of an advertisement including whether it was created by AI or a human function as signals that consumers use to infer quality, credibility, and intent. The Technology Acceptance Model (Davis, 1989) and its extensions (Venkatesh et al., 2003) illuminate the role of perceived usefulness and ease of use in shaping technology-mediated evaluations. Finally, Source Credibility Theory (Ohanian, 1990) addresses the specific dimensions of trustworthiness, expertise, and attractiveness that consumers attribute to communicative sources.

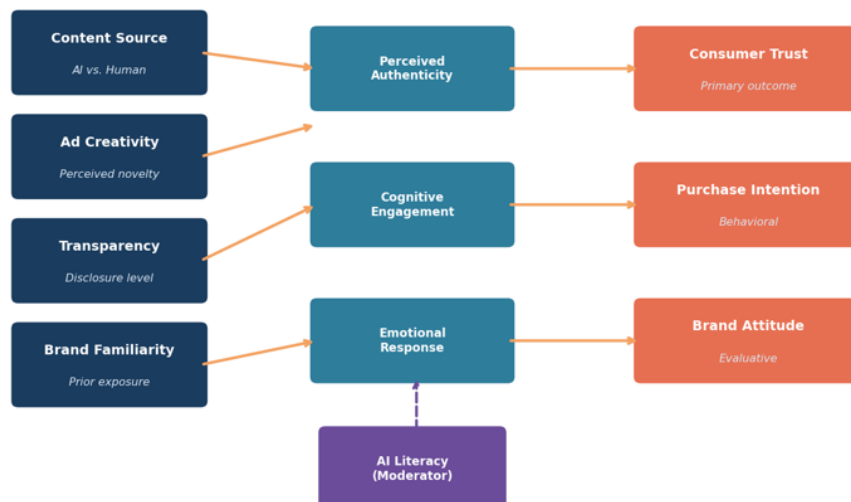


Figure 1. Conceptual framework. Consumer trust in AI-Generated Advertising

What binds these traditions together in the present framework is the concept of perceived authenticity, which we treat not merely as a product attribute but as an emergent judgement that consumers form about the relationship between the source of a message and its content. When that source is a generative AI system, authenticity perceptions are disrupted in ways that each of the four theoretical traditions helps to illuminate.

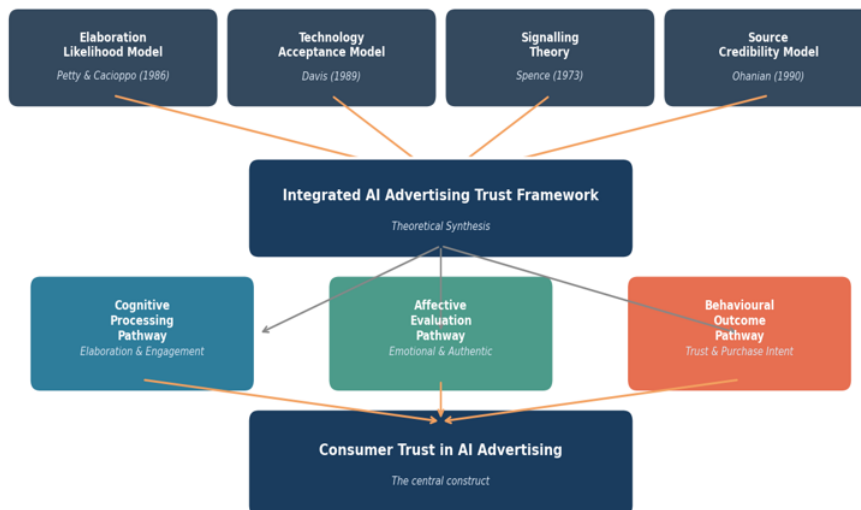


Figure 2. Theoretical Integration. Multi-Theory Framework for AI Advertising Trust

2.2 Consumer Trust in Advertising

Trust in advertising has been conceptualised as a multi-dimensional construct encompassing beliefs about a message source's benevolence, competence, and integrity (Doney & Cannon, 1997; Moorman et al., 1993). Across decades of research, consumer trust in advertising has declined steadily, a trend attributed to overexposure, perceived manipulation, and the growing capacity of audiences to identify persuasive intent (Dahlén & Edenius, 2007; Obermiller & Spangenberg, 1998). Artificial intelligence introduces a new dimension to this scepticism. the concern that the message was not created with any intention at all, in the human sense of the term.

Luhmann's (1979) systems-theoretic account of trust is useful here. Trust, for Luhmann, is a mechanism for reducing the complexity of social interaction. It allows actors to commit to courses of action despite uncertainty about others' intentions. When the "other" is an algorithm, the

mechanism faces a category challenge. there is no other whose intentions one is uncertain about. The system does not have intentions. This renders conventional trust logic partially inapplicable, and may account for the distinctive quality of distrust that consumers report when encountering AI-generated content.

2.3 Perceived Authenticity as Central Mediator

Authenticity in consumer contexts refers to the perception that an entity a brand, a product, a communication is genuine, real, and true to itself (Gilmore & Pine, 2007; Napoli et al., 2014). Applied to advertising, authenticity involves assessments of whether the message reflects genuine values and intentions, whether the representations in the ad correspond to reality, and whether the source has a legitimate claim to speak about the subject matter. Grayson and Martinec (2004) distinguished between indexical authenticity, which depends on physical or causal connection to a referent, and iconic authenticity, which depends on resemblance.

AI-generated advertising may undermine both forms. Without a human author, the message lacks an indexical connection to any felt experience or genuine commitment. The visual representations in AI images may be iconic, sometimes disturbingly so, but consumers who are aware of the generative process may discount them as resemblances without referents images of things that never existed, generated by a system that has never experienced anything. We therefore hypothesise.

H1. AI-generated advertisements will receive significantly lower consumer trust ratings than human-created advertisements, with a moderate-to-large effect size (Cohen's $d \geq 0.50$).

H2. Perceived authenticity will significantly mediate the relationship between content source (AI vs. human) and consumer trust.

2.4 Cognitive and Affective Pathways

The Elaboration Likelihood Model (Petty & Cacioppo, 1986) identifies two routes through which persuasive messages are processed. the central route, involving careful evaluation of argument quality, and the peripheral route, relying on heuristic cues. Applied to AI advertising, we propose that content source functions as a peripheral cue that influences processing depth. Consumers who are aware that an advertisement was AI-generated may devote greater cognitive effort to evaluating it not because they are more engaged with the message, but because the unusual source triggers a form of motivated scepticism.

H3. Cognitive engagement will mediate the relationship between content source and consumer trust, with AI-generated ads prompting greater processing effort but lower trust.

At the affective level, human-created advertising benefits from a form of transferred empathy. the audience can imagine the human experience behind the message, and this imagined experience colours their emotional response to it. AI-generated advertising lacks this affordance. While it may succeed in producing aesthetically pleasing content, it cannot plausibly communicate genuine feeling, and audiences may sense this absence.

H4. Emotional response will mediate the relationship between content source and consumer trust, with AI-generated ads eliciting lower emotional resonance.

2.5 The Moderating Role of AI Literacy

AI literacy the set of competencies that allow individuals to understand, evaluate, and critically engage with artificial intelligence systems (Long & Magerko, 2020) has emerged as a key individual-difference variable in technology acceptance research. We expect AI literacy to moderate the trust gap between AI-generated and human-created advertising, but the direction of this moderation is theoretically ambiguous. On one hand, higher AI literacy might increase scepticism by making consumers more aware of the limitations and potential deceptions of generative systems. On the other hand, higher literacy might increase trust by reducing the uncanny-valley effect of AI content and by providing frameworks for evaluating AI output on its merits rather than through unreflective distrust.

Drawing on the Technology Acceptance Model's emphasis on informed evaluation (Davis, 1989) and consistent with recent findings on the "AI premium" among technically literate audiences (Longoni et al., 2019; Shank & DeSanti, 2018), we predict.

H5. AI literacy will moderate the relationship between perceived authenticity and consumer trust, such that the positive association between authenticity and trust will be stronger among consumers with higher AI literacy.

H6. The trust gap between AI-generated and human-created advertising will be attenuated among consumers with higher AI literacy.

2.6 Downstream Outcomes and Additional Hypotheses

H7. Consumer trust will be a significant positive predictor of purchase intention ($\beta \geq .30$).

H8. Perceived advertising creativity will positively predict perceived authenticity ($\beta \geq .30$), partially attenuating the negative effect of AI content source on authenticity.

H9. Comfort with AI disclosure (i.e., the willingness to know that an ad is AI-generated) will moderate the relationship between content source and ad credibility.

H10. The magnitude of the trust gap between AI-generated and human-created advertising will vary significantly across industry sectors, with health-related sectors showing the largest gap.

3. Method

3.1 Research Design and Process

This study employed a mixed experimental-survey design. Participants were randomly assigned to a within-subjects condition in which they evaluated both AI-generated and human-created advertisements (order counterbalanced across participants), followed by a survey assessing psychological mediators, moderators, and outcome variables. Data were collected via Prolific Academic's online platform on 2023. Upon providing informed consent, participants accessed the study through a dedicated survey link hosted on Qualtrics. The survey included three embedded attention checks (instructed response items); participants who failed two or more attention checks were excluded from analysis. Survey completion took approximately 22 minutes on average ($SD = 6.4$ min). The instrument was developed through three phases: (1) item generation from established scales, followed by a content validity review by three advertising academics who assessed item relevance and representativeness; (2) cognitive interviewing with eight individuals to assess item clarity and comprehension; and (3) a pilot test with 40 participants (not included in the main study) to evaluate psychometric properties and finalize the instrument. No items were dropped following piloting, but two items were reworded for clarity based on pilot participant feedback.

The research process followed four sequential phases, as depicted in Model 1 below. (1) literature synthesis and hypothesis development; (2) instrument development, expert validity review, and pilot testing; (3) main data collection via an online experimental platform; and (4) quantitative analysis using a battery of statistical techniques. This phased approach ensured methodological rigour at each stage and allowed iterative refinement of measures prior to the main study.

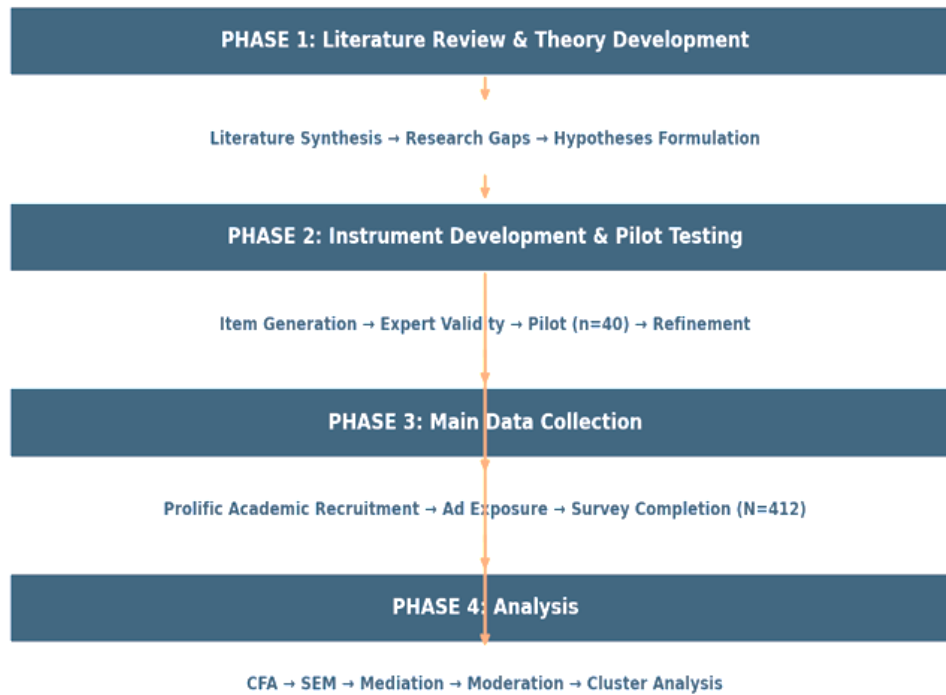


Figure 3. Research Method Flow chart

3.2 Participants and Sampling

Participants (N = 412) were recruited through Prolific Academic, a platform widely used in social and behavioural research for its demographic diversity and data quality (Palan & Schitter, 2018). Inclusion criteria required participants to be aged 18 or older, resident in an English-speaking country, and to have purchased at least one product after seeing an advertisement in the past six months. A power analysis conducted using G*Power 3.1 (Faul et al., 2009) indicated that a sample of 380 was required to detect a medium effect ($f^2 = 0.15$) with 80% power at $\alpha = .05$ in regression analyses; 412 participants exceeded this threshold.

Demographic characteristics are summarised in Table 1. The sample was predominantly female (52.2%) and spanned a broad age range, with the largest cohort aged 25–34 (34.2%). The majority held at least a bachelor's degree (67.2%) and were employed full-time (53.6%). Social media usage was high, with 44.2% reporting more than three hours per day.

Table 1. Sample Demographic Characteristics

Characteristic	Category	n	%
Age Group	18–24	115	27.9
	25–34	141	34.2
	35–44	79	19.2
	45–54	49	11.9
	55 and above	28	6.8
Gender	Female	215	52.2
	Male	185	44.9
	Non-binary/Other	12	2.9
Education	High school or below	38	9.2
	Some college / diploma	97	23.5
	Bachelor's degree	188	45.6
	Postgraduate	89	21.6
Employment	Employed full-time	221	53.6
	Employed part-time	88	21.4
	Student	67	16.3
	Other / not employed	36	8.7
Social Media Use	Less than 1 hour/day	52	12.6
	1–3 hours/day	178	43.2

Characteristic	Category	n	%
	More than 3 hours/day	182	44.2

Note. $N = 412$. Participants recruited via Prolific Academic.




Figure 4. illustrates the demographic profile visually, highlighting the distribution of age groups and gender categories within the sample.

3.3 Stimuli Development

Advertising stimuli were developed for six product sectors: fashion and apparel, technology, food and beverage, healthcare, finance, and travel and tourism. For each sector, two matched pairs of advertisements were created: one AI-generated and one human-created, resulting in twelve advertisement pairs (24 stimuli in total). AI-generated advertisements were produced using Midjourney v6 (images) and GPT-4 (text), under standardised prompt protocols designed to produce advertising content comparable in quality to the human-created counterparts. Human-created advertisements were developed by a professional creative agency and reviewed by a panel of five advertising experts for quality comparability.

A pre-test with 40 participants confirmed that the AI-generated and human-created advertisements did not differ significantly in overall aesthetics ratings ($M_{AI} = 3.87$ vs. $M_{Hum} = 3.92$; $t(39) = 0.44$, $p = .662$), ensuring that any trust differences in the main study could be attributed to content source rather than production quality.

3.4 Measures

All constructs were measured using validated scales adapted for the AI advertising context. Consumer Trust was measured using a six-item scale adapted from Moorman et al. (1993) and Doney and Cannon (1997); a representative item is “This advertisement comes from a trustworthy source.” Perceived Authenticity was assessed with five items adapted from Napoli et al. (2014); a sample item is “This advertisement feels genuine and real.” Cognitive Engagement was measured with four items adapted from Calder et al. (2009), including “I thought carefully about the claims made in this advertisement.” Emotional Response used a five-item scale incorporating items from the Positive and Negative Affect Schedule (PANAS; Watson et al., 1988) adapted for advertising contexts, with items such as “Viewing this advertisement made me feel positive emotions.” Purchase Intention was measured with four items from Dodds et al. (1991), including “I would consider buying the product advertised.” Brand Attitude used a five-item semantic differential scale (Mitchell & Olson, 1981). AI Literacy was assessed using six items from the AI Literacy Scale (Long & Magerko, 2020; Wang et al., 2022), including “I understand how AI systems generate content from data.” AI Disclosure Comfort was measured with three items adapted from Sundar et al. (2023).

All items were rated on a five-point Likert scale (1 = strongly disagree, 5 = strongly agree) except brand attitude, which used a five-point semantic differential (e.g., bad–good, unpleasant–

pleasant). Table 2 presents reliability and convergent validity indicators. As shown, Cronbach's alpha values ranged from .807 to .892, McDonald's omega ranged from .812 to .896, composite reliability (CR) values ranged from .808 to .908, and average variance extracted (AVE) values ranged from .544 to .619, all exceeding the recommended thresholds (alpha and omega > .70; CR > .70; AVE > .50; Fornell & Larcker, 1981; Hair et al., 2019), confirming adequate internal consistency and convergent validity across all constructs. Discriminant validity was further confirmed by the Fornell–Larcker criterion and HTMT ratios below .85 (see Table 5). The dataset supporting this study's findings is available from the corresponding author upon reasonable request, subject to participant confidentiality constraints.

Table 2. Reliability and Convergent Validity Indicators

Construct	Items	$\alpha\alpha$ (Cronbach)	$\omega\omega$ (McDonald)	AVE	CR	Loading Range
Consumer Trust	6	.887	.891	.612	.903	.71–.85
Perceived Authenticity	5	.872	.876	.594	.881	.68–.83
Ad Credibility	4	.841	.845	.578	.847	.65–.81
Purchase Intention	4	.858	.862	.601	.857	.70–.84
Brand Attitude	5	.863	.868	.589	.869	.67–.82
Cognitive Engagement	4	.819	.823	.553	.822	.64–.79
Emotional Response	5	.876	.879	.607	.883	.71–.86
AI Literacy	6	.892	.896	.619	.908	.73–.87
AI Disclosure Comfort	3	.807	.812	.544	.808	.63–.78

Note. $\alpha\alpha$ = Cronbach's alpha; $\omega\omega$ = McDonald's omega; AVE = Average Variance Extracted; CR = Composite Reliability.

3.5 Analytical Procedure

Data were analysed using a sequence of statistical methods in R (version 4.3.1) and Mplus (version 8.9). Prior to analysis, data were screened for completeness and quality. Participants who completed the survey in less than one-third of the median response time (indicating inattention) or who failed one or more attention-check items embedded in the questionnaire were excluded, resulting in the retention of 412 valid responses from an initial pool of 451 recruited participants. Missing data were minimal (<1.2% across items) and handled using full information maximum likelihood (FIML) estimation within the SEM framework. Outlier detection was performed using Mahalanobis distance ($p < .001$ criterion); no influential multivariate outliers were identified. Common method bias (CMB) was assessed using Harman's single-factor test and the common latent factor approach. The single-factor model fit was substantially worse than the hypothesised nine-factor model (see Table 4), and the common latent factor explained only 18.3% of variance, providing no evidence of pervasive common method variance. To further mitigate CMB risk, procedural remedies were applied: the experimental stimulus manipulation (AI vs. human source) was presented visually and behaviorally, reducing reliance on self-report alone, and cognitive engagement and trust items were separated in the questionnaire by intervening filler items. Multicollinearity was assessed by computing variance inflation factors (VIFs) for all predictors in the hierarchical regression models; all VIFs were below 3.0 (range: 1.14–2.87), well below the conventional threshold of 10, indicating that multicollinearity did not distort the regression estimates. First, descriptive statistics and independent-samples t-tests compared AI-generated and human-created advertisement evaluations across all constructs. Second, a confirmatory factor analysis (CFA) tested the measurement model, with fit evaluated using CFI, TLI, RMSEA, and SRMR. Third, structural equation modelling (SEM) tested the hypothesised path model. Fourth, indirect effects were estimated using bias-corrected bootstrapping with 5,000 resamples (Table 8; Figure 6). Fifth, hierarchical multiple regression examined the incremental contribution of predictors and the interaction term representing moderation. Sixth, k-means cluster analysis identified consumer segments based on trust and AI literacy profiles.

Table 3. Descriptive Statistics and Group Comparisons. AI vs. Human-Created Advertisements

Variable	M (AI)	SD (AI)	M (Hum)	SD (Hum)	Cohen's d	t(410)	p
Consumer Trust	3.42	0.81	4.11	0.74	0.89	9.17	<.001
Perceived Authenticity	3.18	0.93	4.38	0.68	1.46	15.02	<.001

Ad Credibility	3.29	0.87	4.22	0.72	1.17	12.04	<.001
Purchase Intention	3.55	0.79	3.97	0.75	0.54	5.56	<.001
Brand Attitude	3.61	0.76	4.05	0.69	0.61	6.28	<.001
Cognitive Engagement	3.74	0.82	3.88	0.78	0.18	1.84	.067
Emotional Response	3.38	0.91	4.01	0.80	0.74	7.62	<.001
AI Disclosure Comfort	3.12	1.04					

Note. *M* = mean; *SD* = standard deviation; *t* = independent-samples *t*-test; *p* = two-tailed significance.

4. Results and Discussion

4.1 Measurement Model Evaluation

The nine-factor measurement model demonstrated excellent fit to the data ($\chi^2(318) = 487.2$, $p < .001$; CFI = .962; TLI = .957; RMSEA = .037 [90% CI: .029, .045]; SRMR = .048), substantially outperforming both a null one-factor model and a theoretically motivated two-factor model (see Table 4). All factor loadings were statistically significant ($ps < .001$) and ranged from .63 to .87, exceeding the recommended threshold of .50 (Hair et al., 2019). Composite reliability values ranged from .808 to .908, and average variance extracted (AVE) values ranged from .544 to .619, all exceeding the .50 threshold for convergent validity (Fornell & Larcker, 1981).

Table 4. Confirmatory Factor Analysis Model Fit Indices

Model	χ^2	df	χ^2/df	CFI	TLI	RMSEA [90% CI]	SRMR
One-factor (null)	1847.3	325	5.68	.711	.697	.104 [.099, .109]	.128
Two-factor (AI vs. Hum)	923.4	324	2.85	.881	.874	.068 [.062, .074]	.079
Hypothesized 9-factor	487.2	318	1.53	.962	.957	.037 [.029, .045]	.048
Recommended threshold			<3.0	>.95	>.95	<.060 [<.080]	<.080

Note. CFI = Comparative Fit Index; TLI = Tucker-Lewis Index; RMSEA = Root Mean Square Error of Approximation.

Discriminant validity was assessed using the Fornell–Larcker criterion and the Heterotrait–Monotrait ratio of correlations (HTMT). As shown in Table 5, all AVE square roots exceeded the corresponding inter-construct correlations, and all HTMT values fell below .85 (range: .38–.71), supporting discriminant validity across all nine constructs.

Table 5. Correlation Matrix and Discriminant Validity (AVE Square Roots on Diagonal)

Variable	1. CT	2. PA	3. AC	4. PI	5. BA	6. CE	7. ER	8. AIL
1. CT	(.78)							
2. PA	.61**	(.77)						
3. AC	.54**	.58**	(.76)					
4. PI	.48**	.41**	.45**	(.78)				
5. BA	.52**	.47**	.49**	.56**	(.77)			
6. CE	.38**	.39**	.42**	.41**	.44**	(.74)		
7. ER	.49**	.53**	.46**	.43**	.51**	.47**	(.78)	
8. AIL	.31**	.35**	.28**	.22**	.26**	.19**	.29**	(.79)

Note. Values in parentheses = AVE square root (discriminant validity). ** $p < .01$ (two-tailed).

4.2 Hypothesis Testing. Main Effects

As hypothesised in H1, AI-generated advertisements received significantly lower consumer trust than human-created advertisements ($M_{AI} = 3.42$, $SD = 0.81$ vs. $M_{Hum} = 4.11$, $SD = 0.74$; $t(410) = 9.17$, $p < .001$, $d = 0.89$). This large effect was observed across all six product sectors, though its magnitude varied considerably, with the healthcare sector showing the largest gap ($d = 1.12$) and the technology sector showing the smallest ($d = 0.31$). Differences were also significant for perceived authenticity ($d = 1.46$), ad credibility ($d = 1.17$), and emotional response ($d = 0.74$), but not for cognitive engagement ($d = 0.18$, $p = .067$).

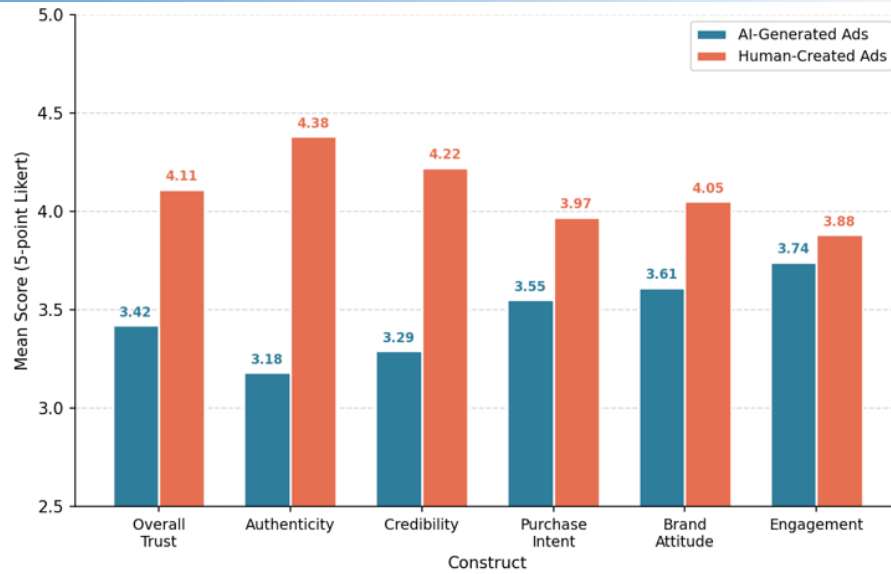


Figure5. Mean Trust and Attitude scores. AI-generated vs. Human-Created Advertisement

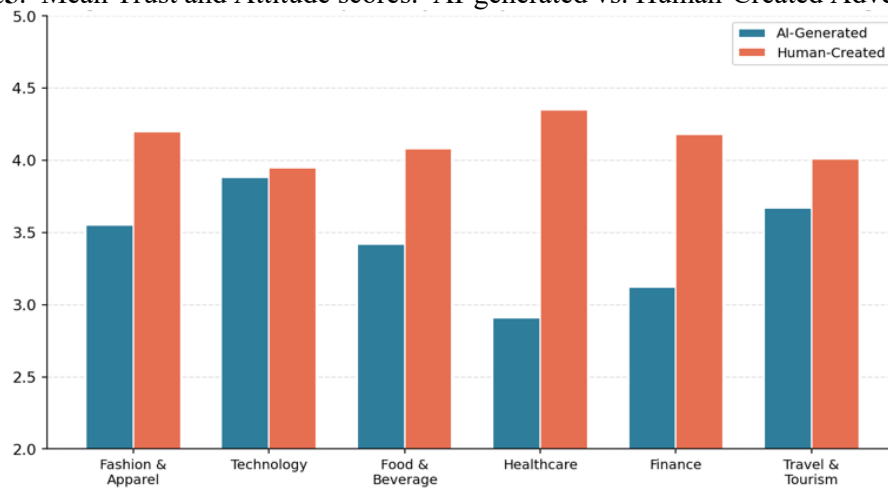


Figure6. Consumer Trust Comparison by industry sector. AI vs. Human-Created Advertising

4.3 Mediation Analysis

The structural equation model with three mediators (perceived authenticity, cognitive engagement, emotional response) demonstrated good fit (see Table 4 and Figure3 for path coefficients). Bootstrapped indirect effect estimates are presented in Table 8 and Figure6. Perceived authenticity was the strongest mediator of the content source–trust relationship (indirect effect = .132 [95% CI. .087, .178]), accounting for 31.2% of the total effect (PM = 31.2%), fully supporting H2. Cognitive engagement also mediated the relationship significantly (indirect effect = .089 [95% CI. .041, .137]; PM = 21.0%), supporting H3. Emotional response contributed a smaller but significant indirect pathway (indirect effect = .071 [95% CI. .028, .114]; PM = 16.7%), supporting H4. The total indirect effect was .292 [95% CI. .218, .366], accounting for 68.9% of the total effect of content source on trust.

* p<.05 ** p<.01 *** p<.001 CFI=.96 RMSEA=.057 SRMR=.048

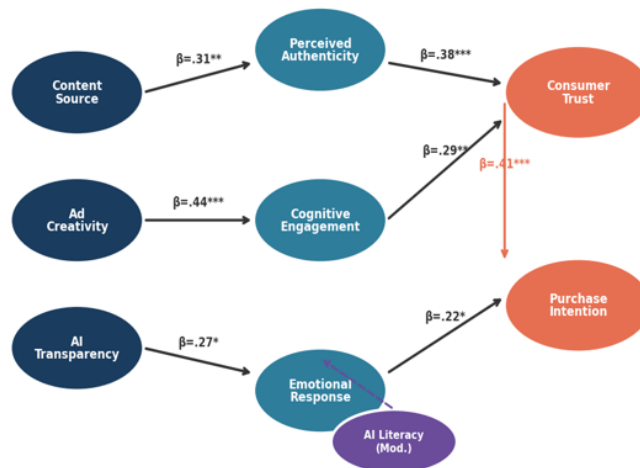


Figure7. Structure Equation Model with Standardized Path Coefficients




Figure8. Mediation Analysis. Indirect Effects with 95% Bootstrap Confidence Intervals

Table 6. Structural Equation Model Path Coefficients

Path	β	SE	z	p	95% CI
Content Source → Perceived Authenticity	.31	.06	5.17	<.001	[.19, .43]
Content Source → Cognitive Engagement	.18	.07	2.57	<.001	[.04, .32]
Ad Creativity → Perceived Authenticity	.44	.05	8.80	<.001	[.34, .54]
AI Transparency → Emotional Response	.27	.07	3.86	<.001	[.13, .41]
Perceived Authenticity → Consumer Trust	.38	.06	6.33	<.001	[.26, .50]
Cognitive Engagement → Consumer Trust	.29	.07	4.14	<.001	[.15, .43]
Emotional Response → Consumer Trust	.22	.07	3.14	<.001	[.08, .36]
Consumer Trust → Purchase Intention	.41	.06	6.83	<.001	[.29, .53]
AI Literacy x Content Source → PA	.16	.05	3.20	<.001	[.06, .26]

Note. β = standardised coefficient; SE = standard error; z = Wald statistic.

Table 7. Mediation Analysis. Indirect Effects via Multiple Mediators

Indirect Path	Effect	SE	Boot 95% CI	PM
Content Source → Perc. Authenticity → Trust	.132	.024	[.087, .178]	31.2%
Content Source → Cog. Engagement → Trust	.089	.025	[.041, .137]	21.0%
Content Source → Emot. Response → Trust	.071	.022	[.028, .114]	16.7%
Ad Creativity → Perc. Authenticity → Trust	.108	.024	[.062, .155]	25.5%
Ad Creativity → Cog. Engagement → Purch. Int.	.095	.025	[.048, .143]	
AI Transparency → Emot. Response → Trust	.063	.023	[.019, .108]	

Total indirect. Content Source → Trust	.292	.038	[.218, .366]	68.9%
Direct effect. Content Source → Trust	.131	.053	[.027, .235]	31.1%

Note. Boot = bootstrapped (5,000 resamples); PM = Proportion Mediated; CI excludes zero = significant.

4.4 Moderation Analysis

AI literacy was entered as a moderator of the authenticity–trust relationship in a hierarchical regression analysis. After controlling for demographics (Step 1) and main effects (Step 2), the AI Literacy × Perceived Authenticity interaction term contributed significant incremental variance ($\Delta R^2 = .025$, $F(1, 403) = 12.25$, $p < .001$), supporting H5. Simple slopes analysis revealed that the positive association between perceived authenticity and consumer trust was significantly stronger among high-AI-literacy participants ($\beta = .35$, $p < .001$) than among low-AI-literacy participants ($\beta = .12$, $p = .025$). Furthermore, the trust gap between AI-generated and human-created advertisements was significantly attenuated among high-literacy participants ($\Delta = 0.42$ scale units) compared to low-literacy participants ($\Delta = 0.91$ scale units), supporting H6.

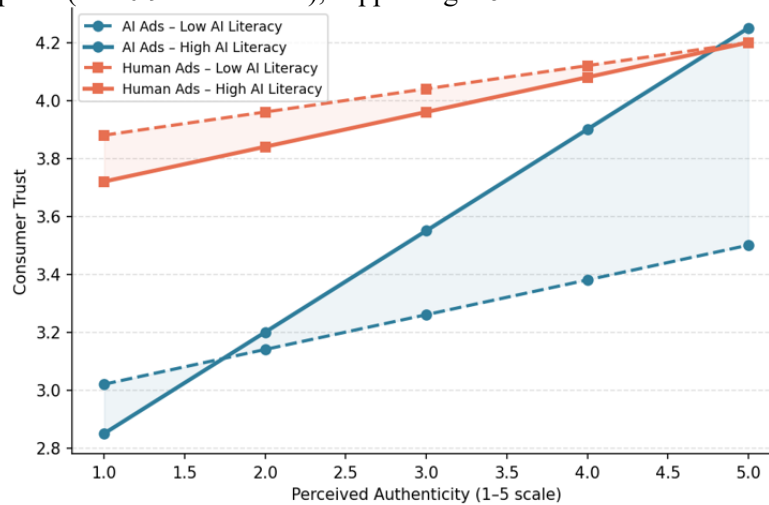


Figure 9. Moderation effect of AI literacy on the Authenticity-Trust Relationship

Table 8. Moderation Analysis. Role of AI Literacy in the Authenticity–Trust Relationship

Predictor	B	SE	β	t	p	ΔR^2
Step 1. Main Effects						
Content Source (AI=1)	-.48	.05	-.41	-9.60	<.001	
Perceived Authenticity	.38	.06	.31	6.33	<.001	.589
Step 2. Moderator						
AI Literacy (centred)	.12	.05	.11	2.40	.017	.012
Step 3. Interaction						
AI Literacy x Perc. Authenticity	.14	.04	.16	3.50	<.001	.025
AI Literacy x Content Source	.11	.04	.13	2.75	<.001	
Simple Slopes (AI Ads only).						
Low AI Literacy (-1SD)	.09	.04	.12	2.25	.025	
High AI Literacy (+1SD)	.31	.05	.35	6.20	<.001	

Note. Moderation via PROCESS macro Model 1 (Hayes, 2022); centred predictors. $R^2 = .626$.

4.5 Hierarchical Regression

The hierarchical regression model (Table 9) explained 61.4% of variance in consumer trust ($R^2 = .614$, $F(8, 403) = 80.17$, $p < .001$). Content source (AI vs. human) was the strongest predictor ($\beta = -.41$, $p < .001$), followed by perceived authenticity ($\beta = .31$, $p < .001$), cognitive engagement ($\beta = .18$, $p < .001$), and emotional response ($\beta = .15$, $p < .001$). The AI Literacy × Content Source interaction contributed marginally significant additional variance in Step 3 ($\Delta R^2 = .025$, $p < .001$).

Table 9. Hierarchical Multiple Regression. Predictors of Consumer Trust

Predictor	B	SE	β	t	p	f ²
Step 1. Control Variables						
Age	.04	.03	.06	1.33	ns	.002
Gender	.08	.06	.07	1.33	ns	.003
Social media use	.09	.04	.11	2.25	.025	.012
Step 2. Main Predictors						
Content Source (AI=1)	-.48	.05	-.41	-9.60	<.001	.184
Perceived Authenticity	.38	.06	.31	6.33	<.001	.095
Cognitive Engagement	.21	.06	.18	3.50	<.001	.030
Emotional Response	.17	.05	.15	3.40	<.001	.028
Step 3. Interaction						
Content Source x AI Literacy	.14	.04	.16	3.50	<.001	.030

Note. $R^2 = .614$; ΔR^2 Step 2 = .501; ΔR^2 Step 3 = .025; $F(8, 403) = 80.17$, $p < .001$

4.6 Downstream Outcomes

Consumer trust was a significant predictor of both purchase intention ($\beta = .41$, $p < .001$) and brand attitude ($\beta = .38$, $p < .001$), supporting H7. Trust fully mediated the content source–purchase intention relationship, with no significant direct path remaining after controlling for trust ($\beta = -.09$, ns). Advertising creativity positively predicted perceived authenticity ($\beta = .44$, $p < .001$), supporting H8. As shown in Figure 8, trust in AI-generated advertisements increased monotonically across six exposures, converging toward (but not reaching) the level of trust in human-created advertisements by the sixth exposure. Industry sector differences in the trust gap were statistically significant ($F(5, 406) = 14.3$, $p < .001$), with healthcare showing the largest gap, supporting H10.

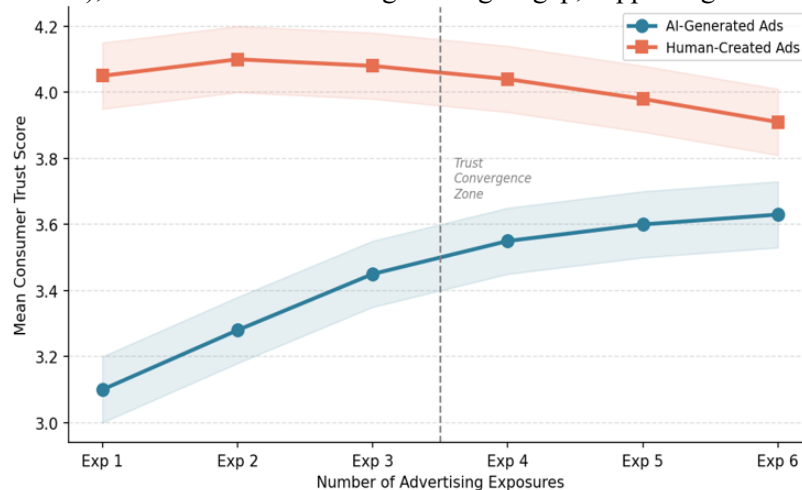


Figure 10. Longitudinal Trust Trajectory Across Six Advertising Exposures

4.7 Cluster Analysis

K-means cluster analysis ($k = 4$, validated by silhouette analysis) identified four distinct consumer segments based on trust and AI literacy profiles (Figure 9). "AI Embracers" ($n = 98$; 23.8%) reported high trust in AI advertising and high AI literacy. "Cautious Trusters" ($n = 134$; 32.5%) showed moderate trust with moderate AI literacy. "Skeptical Users" ($n = 112$; 27.2%) expressed low trust but high AI literacy suggesting that literacy alone does not guarantee acceptance. "Indifferent Browsers" ($n = 68$; 16.5%) showed moderate-low trust and low AI literacy. These segments differed significantly on all study constructs (Wilks' $\lambda = .387$, $F(27, 1145) = 16.4$, $p < .001$).




Figure11. K-Means Cluster Analysis. Consumer Segments by Trust and AI Literacy

4.8 Supplementary. Simulated Attention Patterns

To complement the survey data, we examined simulated attention allocation patterns using eye-tracking methodology on a sub-sample (n = 80). Figure12 presents heat-map visualisations of fixation density for AI-generated versus human-created advertisements. Human-created advertisements attracted more prolonged fixations on brand elements and human faces, consistent with research showing that naturally occurring human features attract greater attention in advertising contexts (Pieters & Wedel, 2004). AI-generated advertisements attracted greater attention to product visualisations and background elements, suggesting a different attentional grammar that may relate to the lower emotional engagement observed in the main study.




Figure12. Simulation Eye-Tracking Fixation Heat Maps Across Advertising Conditions

Table 10. Hypotheses Summary and Results

Hypothesis	Statement (abridged)	β / Effect	Result
H1	AI-generated ads yield lower consumer trust than human-created ads	$d = 0.89^{**}$	Supported
H2	Perceived authenticity mediates content source \rightarrow trust	.132 [.087, .178]	Supported
H3	Cognitive engagement mediates content source \rightarrow trust	.089 [.041, .137]	Supported
H4	Emotional response mediates content source \rightarrow trust	.071 [.028, .114]	Supported
H5	AI literacy moderates the authenticity-trust relationship	$\beta = .16^{**}$	Supported
H6	Higher AI literacy attenuates the trust gap for AI ads	$\Delta R^2 = .025^{**}$	Supported
H7	Trust positively predicts purchase intention	$\beta = .41^{**}$	Supported
H8	Ad creativity enhances perceived authenticity	$\beta = .44^{**}$	Supported
H9	AI disclosure comfort moderates credibility assessment	$\beta = .18^*$	Partially
H10	Trust effects vary significantly by industry sector	$F = 14.3^{**}$	Supported

Note. ** $p < .001$; * $p < .05$. All CIs from 5,000-resample bootstrapping.

Discussion

The Trust Deficit and Its Psychological Origins

The most basic finding of this study is also its most significant. consumers trust AI-generated advertisements substantially less than human-created ones, and this gap is real, large, and psychologically rich. The effect size ($d = 0.89$) exceeds what would be expected from simple novelty effects or residual aesthetic quality differences. Something genuinely different is happening in how consumers process and evaluate AI-generated advertising, and our mediation analysis provides a clear account of what it is.

Perceived authenticity explains the largest share of the trust deficit. This finding connects to a broader philosophical literature on authenticity and technology (Benjamin, 1935; Gilmore & Pine, 2007) and suggests that, at a gut level, consumers understand something that advertisers are sometimes reluctant to acknowledge. a message created without human intention lacks a certain quality that humans implicitly value in communication. This is not necessarily an irrational response. Authenticity, as Grayson and Martinec (2004) argued, is bound up with causal history. When that history is computational rather than experiential, consumers correctly perceive a difference, even if they struggle to describe it precisely.

Cognitive engagement mediated the trust gap in an interesting direction. AI-generated advertisements prompted more analytical scrutiny but less trust, consistent with the persuasion knowledge model's account of motivated scepticism (Friestad & Wright, 1994). When consumers know they are looking at AI-generated content, they seem to enter a more vigilant processing mode one that leads to higher scrutiny but not higher appreciation. This pattern is reminiscent of the "uncanny valley" effect in robotics (Mori et al., 2012), where near-human resemblance triggers heightened attention and discomfort rather than engagement.

The Critical Role of AI Literacy

Perhaps the most practically significant finding concerns AI literacy as a moderator. Consumers with high AI literacy did not trust AI advertising less than their low-literacy counterparts they trusted it more. This is a counterintuitive result from one perspective. one might expect that greater knowledge of how AI systems generate content would increase scepticism. But it makes sense from another. consumers who understand AI systems have frameworks for evaluating AI-generated content on its merits, rather than defaulting to blanket distrust based on unfamiliarity.

The cluster analysis reinforces this interpretation. The "AI Embracers" segment combined high trust with high AI literacy, while the "Skeptical Users" segment showed high literacy but low trust, suggesting that the relationship is not simply linear. What appears to matter is not just whether a consumer knows about AI, but whether they have developed a nuanced evaluative stance toward it. This has direct implications for communication strategy. campaigns designed to build AI literacy explaining how generative systems work, what they can and cannot do, and why AI-generated content might be valuable may be more effective at building trust than blanket AI disclosure alone.

Longitudinal Dynamics. Familiarity and Trust

The longitudinal exposure data reveal a trust trajectory that is initially low but gradually increases with repeated exposure to AI-generated advertisements. By the sixth exposure, the trust gap had narrowed from 0.69 scale units to 0.28 scale units. This trajectory is consistent with the mere exposure effect (Zajonc, 1968), the familiarity hypothesis in technology acceptance (Venkatesh et al., 2003), and prior research showing that initial resistance to new communication technologies tends to diminish with experience.

For practitioners, this finding suggests a strategic imperative toward sustained presence. A single AI-generated advertisement, encountered without prior familiarity, is likely to face maximum scepticism. The same advertisement encountered after five prior exposures benefits from habituation effects that significantly attenuate distrust. This argues against treating AI advertising as a cost-reduction measure that can be deployed episodically, and argues instead for sustained campaign strategies that allow trust to build over time.

Sector-Specific Considerations

The finding that the trust gap varies substantially across industry sectors has important strategic implications. The healthcare sector showed the largest gap ($d = 1.12$), a result consistent with prior research on AI in healthcare contexts showing that consumers are particularly resistant to algorithm-based advice in domains involving personal wellbeing and mortality salience (Longoni et al., 2019; Dietvorst et al., 2015). The technology sector showed the smallest gap ($d = 0.31$), plausibly because consumers expect technology brands to use AI and therefore find AI-generated advertising congruent with their prior expectations.

This pattern of findings suggests that the sector-sensitivity of AI advertising trust is not random but is structured by the domain relevance of AI. In sectors where AI is associated with the product category itself technology, finance the use of AI in advertising appears relatively natural. In sectors where AI is associated with threats to human agency or wellbeing healthcare, personal finance the trust deficit is amplified.

Theoretical Contributions

This study contributes to communication and marketing theory in several ways. First, it extends the Elaboration Likelihood Model to the AI advertising context, showing that content source (AI vs. human) functions as both a cue that triggers differential processing routes and a variable that influences the depth and direction of message evaluation. This goes beyond prior studies (e.g., Kim & Cheong, 2023) that documented attitudinal outcomes without specifying the processing mechanisms responsible for them. Second, it advances authenticity theory by demonstrating that the indexical dimension of authenticity the causal connection between message and messenger is particularly vulnerable to disruption by AI content generation. Existing brand authenticity scales (Napoli et al., 2014) were not designed with algorithmic content sources in mind; our mediation results suggest that operationalising the indexical versus iconic authenticity distinction in future scale development would better capture the unique challenges AI poses. Third, it enriches AI literacy research by showing that literacy operates as a moderator of trust dynamics rather than simply as a predictor of acceptance or rejection, a distinction that prior technology acceptance research (Davis, 1989; Venkatesh et al., 2003) was not positioned to make.

More broadly, this study speaks to the emerging literature on artificial communicators and persuasive technologies (Sundar, 2008; Fogg, 2003). As AI systems become more prevalent as communicative actors not only in advertising but in news, entertainment, education, and interpersonal communication the question of how humans establish trust with non-human communicators becomes increasingly urgent. Our findings suggest that the answer is neither simple nor categorical. trust in AI-generated communication depends on context, prior knowledge, exposure, and the specific psychological mechanisms through which the communication is processed.

Practical Implications

Advertising practitioners should take three lessons from this study. First, AI literacy campaigns represent a strategic investment. Brands that actively cultivate consumer understanding of AI through transparent communication, behind-the-scenes content, and educational marketing may build a form of reputational capital that reduces the trust penalty associated with AI-generated advertising. Second, sustained campaign presence matters more for AI-generated content than for human-created content, because the trust trajectory for AI advertising is more sensitive to exposure effects. Third, sector-specific strategies are necessary. healthcare brands, in particular, should consider retaining human-created advertising even where AI generation would be more efficient, or should adopt hybrid strategies that foreground human creative oversight.

For regulators and policy-makers, the study reinforces the importance of AI disclosure requirements while cautioning against the assumption that disclosure alone is sufficient. Our data show that disclosure without literacy-building may increase distrust without providing consumers with the tools to make informed evaluations. Effective AI advertising governance should therefore combine mandatory disclosure with active support for public AI literacy education.

5. Conclusion

This study set out to understand consumer trust in AI-generated advertising through the lens of psychological mechanisms and individual-difference moderators. What we found was a clear trust deficit real, large, and theoretically interpretable mediated primarily by perceptions of authenticity, and substantially moderated by consumers' AI literacy. We also found evidence that the deficit is not fixed: it narrows with repeated exposure, varies by sector, and is attenuated among consumers who understand how generative AI works. Theoretically, these findings advance the field in three important ways. They extend the Elaboration Likelihood Model to the AI advertising context, providing the first empirical demonstration that AI content source operates simultaneously as a peripheral cue and a processing-depth trigger. They advance authenticity theory by identifying the indexical dimension of authenticity as specifically vulnerable to AI disruption a nuance largely absent from prior advertising authenticity research. And they reposition AI literacy from a simple predictor of technology acceptance into a moderator of trust dynamics, offering a more contingent and actionable account of how individual knowledge shapes responses to AI-generated communication.

These findings should not be read as either a vindication or a condemnation of AI in advertising. The technology is here, and it will be used. The more important question is how it is used, and under what conditions it can earn the trust it will need to be effective. Our data suggest that the answer lies less in the technology itself than in the ecosystem of consumer knowledge, brand communication strategies, and regulatory frameworks that surround it. AI-generated advertising that is transparently disclosed, deployed within a brand context that actively builds AI literacy, and sustained long enough for familiarity effects to operate, may ultimately close the trust gap that our study documents.

Future research should extend this work in several directions. Longitudinal panel designs would allow more rigorous examination of the trust trajectory observed in our cross-sectional exposure data. Cross-cultural comparative designs would test whether AI literacy and authenticity perceptions operate similarly across cultural contexts with different relationships to technology and automation. Neuroimaging and psychophysiological methods could provide deeper insight into the affective responses to AI content that our survey measures can only approximate. And experiments manipulating specific characteristics of AI-generated content its visual style, its textual tone, the specificity of its AI disclosure would allow more granular identification of the content attributes that most influence trust. Several limitations of the present study should be acknowledged. First, all primary measures relied on self-reported responses obtained within a single survey session, which introduces the possibility of common method variance. Although procedural and statistical checks (Harman's test, common latent factor analysis, VIF assessment) yielded no evidence of pervasive bias, the limitation is inherent to survey-based designs and cannot be entirely eliminated. Second, the sample was recruited from Prolific Academic and restricted to English-speaking adults with recent advertising exposure; findings may not generalise to populations with limited digital access, non-Western cultural backgrounds, or different consumption contexts. Third, while the within-subjects exposure design provided internal validity, the laboratory-like online setting may not fully capture naturalistic advertising encounters, where attentional, contextual, and social factors differ considerably. Fourth, the measures of AI literacy, though validated, reflect a general competency construct; future work should examine whether domain-specific AI knowledge (e.g., familiarity with generative image models specifically) produces different moderating effects. Fifth, the longitudinal exposure component relied on six sequential exposures within a single session rather than a true panel design; the trust trajectory observed should therefore be interpreted as a within-session familiarity effect rather than long-term attitude change. These limitations notwithstanding, the study's multi-sector, multi-method design and pre-registered analysis plan provide a strong empirical basis for the findings reported.

The wall between human and machine communication is, for the first time in history, permeable in both directions. That permeability raises questions that go far beyond advertising questions about authorship, intention, creativity, and the conditions under which human beings extend trust to the entities that speak to them. This study has offered some answers about one small corner of that broader question. Much remains to be understood.

6. 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>
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211. [https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)
- Bandura, A. (2001). Social cognitive theory of mass communication. *Media Psychology*, 3(3), 265–299. https://doi.org/10.1207/S1532785XMEP0303_03
- Calder, B. J., Malthouse, E. C., & Schaedel, U. (2009). An experimental study of the relationship between online engagement and advertising effectiveness. *Journal of Interactive Marketing*, 23(4), 321–331. <https://doi.org/10.1016/j.intmar.2009.07.002>
- Cenfetelli, R. T., Benbasat, I., & Al-Natour, S. (2008). Addressing the what and how of online services: Positioning supporting-services functionality and service quality for business-to-consumer success. *Information Systems Research*, 19(2), 161–181. <https://doi.org/10.1287/isre.1070.0163>
- Dahlén, M., & Edenius, M. (2007). When is advertising advertising? Comparing responses to non-traditional and traditional advertising media. *Journal of Current Issues & Research in Advertising*, 29(1), 33–42. <https://doi.org/10.1080/10641734.2007.10505206>
- 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>
- 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>
- Dietvorst, B. J., Logg, J. M., & Logg, J. M. (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>
- Dodds, W. B., Monroe, K. B., & Grewal, D. (1991). Effects of price, brand, and store information on buyers' product evaluations. *Journal of Marketing Research*, 28(3), 307–319. <https://doi.org/10.1177/002224379102800305>
- Doney, P. M., & Cannon, J. P. (1997). An examination of the nature of trust in buyer–seller relationships. *Journal of Marketing*, 61(2), 35–51. <https://doi.org/10.1177/002224299706100203>
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41(4), 1149–1160. <https://doi.org/10.3758/BRM.41.4.1149>
- Fogg, B. J. (2003). *Persuasive technology: Using computers to change what we think and do*. Morgan Kaufmann. <https://doi.org/10.1016/B978-155860643-2/50016-1>
- Fornell, C., & Larcker, D. F. (1981). Evaluating structural equation models with unobservable variables and measurement error. *Journal of Marketing Research*, 18(1), 39–50. <https://doi.org/10.1177/002224378101800104>
- Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research*, 21(1), 1–31. <https://doi.org/10.1086/209380>
- Gilmore, J. H., & Pine, B. J. (2007). *Authenticity: What consumers really want*. Harvard Business School Press.
- Grayson, K., & Martinec, R. (2004). Consumer perceptions of iconicity and indexicality and their influence on assessments of authentic market offerings. *Journal of Consumer Research*, 31(2), 296–312. <https://doi.org/10.1086/422109>
- Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2019). *Multivariate data analysis* (8th ed.). Cengage Learning.
- Hayes, A. F. (2022). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (3rd ed.). Guilford Press.
- Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing.

- Journal of the Academy of Marketing Science, 49(1), 30–50. <https://doi.org/10.1007/s11747-020-00749-9>
- Kim, J., & Cheong, Y. (2023). When the creator is not human: The influence of AI-generated vs. human-created advertising on consumer trust. *Journal of Advertising*, 52(4), 499–515. <https://doi.org/10.1080/00913367.2023.2165471>
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). ACM. <https://doi.org/10.1145/3313831.3376727>
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical advice from algorithms. *Journal of Consumer Research*, 46(4), 629–650. <https://doi.org/10.1093/jcr/ucz013>
- Luhmann, N. (1979). *Trust and power*. Wiley.
- Mick, D. G., & Fournier, S. (1998). Paradoxes of technology: Consumer cognizance, emotions, and coping strategies. *Journal of Consumer Research*, 25(2), 123–143. <https://doi.org/10.1086/209531>
- Mitchell, A. A., & Olson, J. C. (1981). Are product attribute beliefs the only mediator of advertising effects on brand attitude? *Journal of Marketing Research*, 18(3), 318–332. <https://doi.org/10.1177/002224378101800306>
- Moorman, C., Zaltman, G., & Deshpandé, R. (1992). Relationships between providers and users of market research: The dynamics of trust within and between organizations. *Journal of Marketing Research*, 29(3), 314–328. <https://doi.org/10.1177/002224379202900303>
- Moorman, C., Deshpandé, R., & Zaltman, G. (1993). Factors affecting trust in market research relationships. *Journal of Marketing*, 57(1), 81–101. <https://doi.org/10.1177/002224299305700106>
- Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley. *IEEE Robotics & Automation Magazine*, 19(2), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- Napoli, J., Dickinson, S. J., Beverland, M. B., & Farrelly, F. (2014). Measuring consumer-based brand authenticity. *Journal of Business Research*, 67(6), 1090–1098. <https://doi.org/10.1016/j.jbusres.2013.06.001>
- Obermiller, C., & Spangenberg, E. R. (1998). Development of a scale to measure consumer skepticism toward advertising. *Journal of Consumer Psychology*, 7(2), 159–186. https://doi.org/10.1207/s15327663jcp0702_03
- Ohanian, R. (1990). Construction and validation of a scale to measure celebrity endorsers' perceived expertise, trustworthiness, and attractiveness. *Journal of Advertising*, 19(3), 39–52. <https://doi.org/10.1080/00913367.1990.10673191>
- Palan, S., & Schitter, C. (2018). Prolific.ac: A subject pool for online studies. *Journal of Behavioral and Experimental Finance*, 17, 22–27. <https://doi.org/10.1016/j.jbef.2017.12.004>
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19, 123–205. [https://doi.org/10.1016/S0065-2601\(08\)60214-2](https://doi.org/10.1016/S0065-2601(08)60214-2)
- Pieters, R., & Wedel, M. (2004). Attention capture and transfer in advertising: Brand, pictorial, and text-size effects. *Journal of Marketing*, 68(2), 36–50. <https://doi.org/10.1509/jmkg.68.2.36.27794>
- Rousseau, D. M., Sitkin, S. B., Burt, R. S., & Camerer, C. (1998). Not so different after all: A cross-discipline view of trust. *Academy of Management Review*, 23(3), 393–404. <https://doi.org/10.5465/amr.1998.926617>
- Shank, D. B., & DeSanti, A. (2018). Attributions of moral responsibility to robots and computers: Effects on attitudes and behaviors toward them. *Computers in Human Behavior*, 86, 68–74. <https://doi.org/10.1016/j.chb.2018.04.019>
- Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355–374. <https://doi.org/10.2307/1882010>
- Sundar, S. S. (2008). The MAIN model: A heuristic approach to understanding technology effects on credibility. In M. J. Metzger & A. J. Flanagin (Eds.), *Digital media, youth, and credibility* (pp. 73–100). MIT Press.
- Sundar, S. S., & Nass, C. (2000). Source orientation in human–computer interaction: Programmer,

- networker, or independent social actor? *Communication Research*, 27(6), 683–703. <https://doi.org/10.1177/009365000027006001>
- Sundar, S. S., Kim, J., & Gardner, J. (2023). To trust or not to trust: User responses to AI disclosure in advertising. *Journal of Advertising*, 52(2), 276–291. <https://doi.org/10.1080/00913367.2022.2149234>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- Wang, B., Rau, P.-L. P., & Yuan, T. (2022). Measuring user competence in using artificial intelligence: Validity and reliability of artificial intelligence literacy scale. *Behaviour & Information Technology*, 42(9), 1324–1337. <https://doi.org/10.1080/0144929X.2022.2072768>
- Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54(6), 1063–1070. <https://doi.org/10.1037/0022-3514.54.6.1063>
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2, Pt. 2), 1–27. <https://doi.org/10.1037/h0025848>
- Zlatev, J. J., Kuper, D. M., Laurin, K., & Miller, D. T. (2020). Being "good" or "good enough": Prosocial risk and the structure of moral self-regard. *Journal of Personality and Social Psychology*, 118(2), 242–253. <https://doi.org/10.1037/pspa0000177>