

Consumer Skepticism Toward Artificial Intelligence–Generated Brand Narratives

Dr. Olivia Thompson

Associate Professor, University of Manchester, UK

Abstract

Artificial intelligence is increasingly relevant to marketing communication because algorithmic systems can support the production, personalization, and distribution of textual content at considerable scale. Brand narratives, however, differ from routine informational messages because their effectiveness commonly depends on perceived authenticity, emotional resonance, narrative coherence, symbolic meaning, and confidence that the brand is communicating truthfully. These characteristics create a potential tension when consumers learn that a narrative has been generated substantially by artificial intelligence.

The present study examines consumer skepticism toward artificial intelligence–generated brand narratives through a theoretically grounded simulation incorporating perceived narrative authenticity, narrative coherence, disclosure trust, human oversight, and consumer skepticism. Because direct empirical research specifically examining contemporary generative-AI brand narratives was not mature before 2021, the conceptual framework integrates established pre-2021 evidence from advertising skepticism, brand authenticity, narrative transportation, automated journalism, algorithm aversion, and artificial intelligence in marketing. A synthetic analytical corpus of 160 consumer evaluations was constructed across four narrative-source conditions: human-authored, disclosed AI-assisted, disclosed AI-generated with human review, and AI-generated without human review.

The simulation produced a negative association of $r = -0.732$ between perceived narrative authenticity and consumer skepticism. Mean skepticism increased from 36.2 for human-authored narratives to 43.7 for disclosed AI-assisted narratives, 53.6 for AI-generated narratives with human review, and 63.0 for AI-generated narratives without human review. The findings indicate that skepticism should not be interpreted as an automatic rejection of machine-produced language. Rather, consumers may evaluate whether the narrative appears truthful, coherent with brand history, appropriately disclosed, and subject to accountable human oversight.

The study proposes a responsible brand-narrative framework in which artificial intelligence operates as a creative support mechanism while factual claims, value statements, testimonials, cultural representations, and emotionally consequential storytelling remain subject to human verification and organizational accountability.

Keywords: artificial intelligence; brand narratives; consumer skepticism; brand authenticity; AI-generated content; narrative transportation; brand trust; disclosure; algorithm aversion; marketing communication

1. Introduction

Brand narratives are important marketing instruments because stories can organize information into characters, events, motivations, conflicts, and outcomes that consumers can interpret more readily than isolated promotional claims. Research on narrative transportation demonstrates that consumers can become cognitively and emotionally involved in story worlds, and meta-analytic evidence indicates that narrative transportation can influence attitudes, intentions, and other persuasive outcomes. Digital storytelling research has further shown that characteristics of the storyteller and communication environment can moderate narrative effects. These findings explain why brands frequently use stories about founders, customers, heritage, communities, product origins, social purpose, innovation, and organizational values to create meaning beyond functional product information.

Authenticity is particularly important in this context. Beverland and Farrelly's consumer research showed that people actively seek cues they interpret as genuine, real, or true when evaluating consumption objects, while Morhart and colleagues developed a multidimensional perceived brand authenticity scale containing continuity, credibility, integrity, and symbolism. A brand narrative can therefore fail even when it is grammatically polished if consumers perceive it as inconsistent with the brand's history, artificially emotional, commercially opportunistic, or disconnected from recognizable organizational behavior. Brand storytelling is consequently not simply a text-generation problem; it is an authenticity-management problem.

Artificial intelligence complicates this relationship because automation can increase marketing scale and analytical capability while altering perceptions of authorship and agency. Davenport and colleagues argued in 2020 that AI was likely to reshape both marketing strategy and customer behavior across different intelligence levels and marketing tasks. Adjacent empirical research on automated journalism already showed before 2021 that audiences did not invariably reject machine-generated text. Graefe and Bohlken's 2020 meta-analysis of 12 studies involving 4,473 participants found no overall credibility disadvantage for automated news, although human-written news had advantages in perceived quality and especially readability; experimental labeling also influenced evaluations. Clerwall's earlier automated-news experiment likewise found relatively limited differences between software- and journalist-produced factual news, demonstrating that machine authorship does not necessarily produce universal skepticism.

Brand narratives may nevertheless present a different psychological problem from routine factual reports. Research on algorithm aversion shows that people can become disproportionately reluctant to rely on algorithms after observing them make mistakes, while later marketing research demonstrated that algorithm aversion varies according to the nature of the task. This task-dependent perspective is important because brand storytelling frequently involves emotion, identity, values, social meaning, cultural interpretation, and representations of lived experience. It is therefore reasonable to hypothesize—not to assume as an established pre-2021 empirical fact—that consumers may apply

stronger authenticity expectations to AI-generated brand stories than to automated factual summaries. The present paper treats this proposition as a research question requiring empirical validation rather than as an already proven conclusion.

The study consequently examines Consumer Skepticism Toward Artificial Intelligence–Generated Brand Narratives through a simulation-based framework that separates source condition from message quality. Skepticism is conceptualized as a consumer's tendency to question the truthfulness, sincerity, authenticity, or persuasive intent of a brand narrative. This definition draws on Obermiller and Spangenberg's established conceptualization of skepticism toward advertising while adapting the construct to contemporary algorithmically produced storytelling. The central argument is that AI authorship need not automatically undermine brand communication when its role is disclosed appropriately, the narrative remains coherent with verifiable brand identity, and humans retain accountable oversight. Conversely, opacity, unsupported emotional claims, synthetic testimonials, fabricated experiences, or poorly supervised automated storytelling may increase skepticism because they weaken the connection between narrative voice and organizational responsibility.

2. Objectives of the Study

2.1 To Examine the Relationship Between AI Authorship and Consumer Skepticism

The first objective is to examine whether different levels of artificial-intelligence involvement in brand-story production may correspond to different levels of consumer skepticism. Rather than defining the research design through a simple human-versus-machine dichotomy, the study distinguishes human-authored narratives, AI-assisted narratives with disclosure, AI-generated narratives with human review, and AI-generated narratives without substantive human review. This distinction is theoretically justified because the perceived agency behind automated content can influence how audiences evaluate a message, and early automated-journalism evidence showed that authorship information itself can affect judgments independently of message content.

The objective therefore concerns the structure of authorship, not merely the presence of AI. Human-AI collaboration may be perceived differently from autonomous generation because human involvement provides a potential point of responsibility for checking factual claims, tone, cultural context, and alignment with brand values. The study consequently evaluates whether consumer skepticism increases most strongly when AI generation is combined with weak human accountability rather than simply whenever AI is involved.

2.2 To Assess Authenticity, Coherence, and Disclosure as Skepticism Mechanisms

The second objective is to examine whether perceived authenticity and narrative coherence provide a more informative explanation of skepticism than AI authorship alone. Brand-authenticity scholarship emphasizes credibility, integrity, continuity, and symbolic consistency, while narrative research demonstrates that story structure and perceived verisimilitude influence engagement. A technically fluent AI-generated story may therefore receive a favorable response when it accurately reflects established brand history and identifiable values, whereas a human-written narrative may still produce skepticism if it appears manipulative or inconsistent.

Disclosure is treated separately from authenticity. Transparent disclosure may initially make algorithmic involvement more salient, but secrecy creates a different risk: consumers who later discover undisclosed automation may reinterpret the narrative as deceptive. The study therefore treats disclosure as an accountability mechanism rather than assuming that disclosure necessarily increases or decreases persuasion in every situation. Early automated-content research showed that authorship labels can affect evaluations, reinforcing the need to test disclosure rather than treat it as behaviorally neutral.

2.3 To Develop a Responsible AI Brand-Narrative Governance Framework

The third objective is to develop a governance framework that permits brands to use artificial intelligence without treating consumer skepticism as an obstacle to be manipulated away. The proposed approach emphasizes truthful disclosure where material, verification of factual claims, preservation of brand continuity, human review of sensitive narratives, and prohibition of fabricated personal experiences presented as genuine testimony.

This approach follows the broader pre-2021 marketing literature suggesting that AI should be understood according to the tasks it performs and the human–technology configuration surrounding those tasks. It also reflects authenticity research showing that consumer judgments are linked to credibility and integrity rather than surface polish alone. The relevant strategic question is consequently not whether AI can produce convincing stories, but whether brands can use AI in ways that preserve truthful representation, accountable authorship, and credible relationships with consumers.

3. Review of Literature

3.1 Consumer Skepticism, Persuasion Knowledge, and Brand Authenticity

Consumer skepticism toward promotional communication predates artificial intelligence. Obermiller and Spangenberg developed a nine-item scale in 1998 to measure consumers' generalized skepticism toward advertising, defining skepticism as a tendency toward disbelief in advertising claims. This literature is valuable for AI-generated brand narratives because it demonstrates that skepticism should not be interpreted simply as technological fear. Consumers already approach persuasive communication with varying levels of doubt, and the introduction of AI adds a new source cue to an existing process of evaluating commercial credibility.

Advertising skepticism can become particularly relevant when the narrative format appears to blur the distinction between persuasion and authentic personal communication. Consumers may ask whether a founder story was actually written by the founder, whether a customer narrative reflects a genuine experience, whether an emotional description represents the organization's authentic voice, or whether apparently intimate language has been generated at scale. These questions do not necessarily imply that AI use is illegitimate. They indicate that the perceived relationship between speaker, experience, and claim matters to the meaning of the narrative.

Brand-authenticity research provides a stronger theoretical vocabulary for this problem. Beverland and Farrelly found that consumers actively use authenticity cues to obtain identity-related benefits from consumption, while Morhart and colleagues conceptualized perceived brand authenticity through continuity, credibility, integrity, and symbolism. AI-generated storytelling can therefore be

evaluated against four practical questions: Does the story remain consistent with the brand's established identity? Is it credible? Does it reflect stated values with integrity? Does it express meanings consumers recognize as genuinely connected to the brand? These dimensions provide a more nuanced foundation for studying skepticism than a simple measure of whether consumers like or dislike artificial intelligence.

3.2 Narrative Transportation and the Importance of the Storyteller

Narrative transportation theory explains why stories can be unusually persuasive. Van Laer, de Ruyter, Visconti, and Wetzels synthesized two decades of research and described narrative transportation as a process in which consumers mentally enter a story world, with identifiable antecedents and consequences for persuasion. Later digital-era meta-analysis by van Laer, Feiereisen, and Visconti examined 64 articles and 138 effect sizes and found that narrative effects vary according to factors including whether the story is user-generated or professionally generated and how it is received. These findings imply that the source of a story is not merely metadata; it can shape the consumer's interpretive relationship with the narrative.

Marketing research also indicates that narrative structure, characters, and verisimilitude contribute to engagement. Dessart's work found that story plots, characters, and perceived realism can activate cognitive, emotional, and behavioral responses. Studies of narrative advertising similarly show that storytelling is not always superior to informational formats; narrative effectiveness depends on the advertising context and the manner in which the story is presented. This qualification is particularly useful for AI-generated brand communication because it prevents an assumption that more emotionally elaborate AI storytelling will necessarily improve outcomes.

For artificial-intelligence-generated narratives, the storyteller becomes conceptually ambiguous. A brand may specify the purpose, provide factual material, define the tone, prompt the system, revise the output, and approve the final text. The machine may generate the surface language, but organizational responsibility remains human. The relevant empirical question is therefore whether consumers distinguish linguistic authorship from organizational authorship. The present study proposes that human review can serve as a credibility signal because it preserves accountable responsibility without falsely suggesting that every sentence originated from a human writer.

3.3 Automated Content, Algorithm Aversion, and the Pre-2021 Evidence Boundary

Pre-2021 automated-journalism research provides valuable adjacent evidence because it directly compared audience reactions to machine- and human-produced texts. Clerwall's 2014 study investigated users' perceptions of automated news, while Jung and colleagues examined public and journalist perceptions of news written by algorithms and humans. Zheng, Zhong, and Yang later investigated automated-news perceptions across cultural contexts, illustrating that responses to algorithmically produced communication can vary across audiences and settings.

The strongest pre-2021 synthesis is Graefe and Bohlken's 2020 meta-analysis. Across 12 studies and 4,473 participants, they reported no overall credibility difference between automated and human-written news, a small quality advantage for human-written news, and a much larger readability

advantage for human writing. Their experimental synthesis also indicated that simply telling participants an article was human-written improved ratings, creating an ethical tension around disclosure. These findings are particularly valuable because they demonstrate that machine authorship does not inherently destroy credibility, while source labels can nevertheless affect perception.

Algorithm-aversion research adds another layer. Dietvorst, Simmons, and Massey found that people may become especially reluctant to use algorithms after observing them make errors, even when algorithms outperform human forecasters. Castelo, Bos, and Lehmann later demonstrated that algorithm aversion is task dependent, which suggests that reactions to AI may differ according to whether consumers perceive a task as objective, mechanical, subjective, emotional, or requiring human judgment. Because emotionally expressive brand narratives involve symbolic and relational judgments, the present study treats higher skepticism toward autonomous AI storytelling as a theoretically plausible hypothesis rather than as a direct finding already established in the pre-2021 literature.

4. Research Methodology

4.1 Research Design and Simulation Framework

The study employs a quantitative simulation-based research design. No real consumer participants, brand campaigns, product categories, AI platforms, advertisements, or survey responses are represented in the numerical findings. Simulation is used because direct pre-2021 empirical evidence specifically examining modern generative-AI brand narratives is limited and because no verified primary dataset was supplied for the present manuscript.

A synthetic analytical corpus of 160 consumer evaluations was constructed using a fixed random seed, with 40 modeled observations in each of four narrative-source conditions. The conditions were Human-Authored, AI-Assisted with Disclosure, AI-Generated with Human Review, and AI-Generated without Human Review. Four standardized variables were modeled on 0–100 scales: Perceived Narrative Authenticity, Narrative Coherence, Disclosure Trust, and Consumer Skepticism.

The analytical structure is intentionally probabilistic rather than deterministic. Consumers in the AI-generated conditions are not automatically assigned high skepticism, and some simulated human-authored narratives receive relatively low authenticity. This design reflects the theoretical position that message evaluation depends on both source information and perceived narrative qualities.

4.2 Construct Operationalization and Proposed Questionnaire

Consumer Skepticism is operationalized as doubt regarding the sincerity, truthfulness, authenticity, and persuasive legitimacy of a brand narrative. The construct adapts established advertising-skepticism research to the AI-generated communication context. Perceived Narrative Authenticity incorporates whether the story feels consistent with the brand's identity, values, history, and recognizable behavior, drawing conceptually on established brand-authenticity research.

For future empirical validation, the following five-item questionnaire is proposed. It has not been administered to actual consumers. A future experiment could use a five-point or seven-point Likert response format after participants read otherwise comparable brand narratives presented under different authorship and disclosure conditions.

Table 1: Proposed Measurement Instrument for AI-Generated Brand Narratives

Item	Construct	Proposed Questionnaire Statement	Measurement Purpose
Q1	Narrative authenticity	This brand narrative feels consistent with what the brand genuinely represents.	Measures perceived authenticity
Q2	Narrative credibility	I believe the factual and experiential claims presented in this narrative are trustworthy.	Measures credibility
Q3	Disclosure trust	The brand is sufficiently transparent about the role of artificial intelligence in producing this narrative.	Measures authorship transparency
Q4	Human accountability	I am more comfortable with AI-generated brand storytelling when a responsible human has reviewed and approved the final narrative.	Measures perceived value of oversight
Q5	Consumer skepticism	I am cautious about accepting this narrative at face value because its persuasive content may not fully reflect genuine human experience or brand intention.	Measures skepticism

Future empirical research should randomize identical or closely matched texts across authorship labels to distinguish message effects from source-label effects. This is particularly important because automated-journalism research demonstrates that readers' assessments can change when the perceived source changes even when underlying content is comparable. Researchers should also measure preexisting attitudes toward AI, general advertising skepticism, brand familiarity, prior brand trust, AI literacy, and narrative transportation as potential covariates.

4.3 Analytical Procedure and Evaluation Strategy

The simulation assigns different baseline authenticity and skepticism expectations to the four narrative-source conditions while preserving substantial within-group variation. Human-authored narratives were given the highest mean authenticity baseline, followed by disclosed AI-assisted narratives, AI-generated narratives with human review, and AI-generated narratives without human review. These assumptions are analytical propositions, not reported empirical estimates.

Consumer Skepticism was modeled as inversely associated with Perceived Narrative Authenticity, Narrative Coherence, and Disclosure Trust while retaining random variation. The purpose was to examine whether authenticity provides an interpretable pathway connecting AI involvement with skepticism. Pearson correlation was calculated descriptively between Perceived Narrative Authenticity and Consumer Skepticism.

The analysis emphasizes effect patterns rather than inferential significance because significance testing on synthetic data would not establish population-level evidence. Future empirical research should use randomized experiments, multivariate regression, mediation analysis, or structural equation modeling to test whether AI-authorship disclosure influences skepticism directly or indirectly through perceived authenticity, credibility, social presence, and trust.

5. Analysis

5.1 Comparative Skepticism Across Narrative-Source Conditions

The synthetic results reveal systematic differences among the four modeled narrative-source conditions. Human-authored narratives produced the lowest mean Consumer Skepticism Index at 36.2, while disclosed AI-assisted narratives produced a moderately higher value of 43.7. AI-generated narratives subject to human review generated a mean skepticism score of 53.6, and AI-generated narratives without substantive human review produced the highest modeled skepticism at 63.0.

These results illustrate a graded relationship rather than a binary human-versus-AI effect. The relatively moderate difference between Human-Authored and AI-Assisted conditions suggests that disclosed AI involvement may be less problematic when consumers still perceive a meaningful human role in creative direction and accountability. The considerably higher skepticism assigned to unsupervised AI generation illustrates the study's central hypothesis that consumers may react not simply to machine involvement but to uncertainty concerning responsibility for the narrative.

Table 2: Simulated Consumer Response Across Brand-Narrative Source Conditions

Narrative-Source Condition	Simulated Cases	Mean Perceived Authenticity	Mean Narrative Coherence	Mean Disclosure Trust	Mean Consumer Skepticism
Human-authored	40	76.9	73.0	68.9	36.2
AI-assisted, disclosed	40	65.1	68.9	71.6	43.7
AI-generated, human-reviewed	40	57.1	67.4	71.2	53.6
AI-generated, no human review	40	47.8	67.6	57.7	63.0

The coherence scores remain relatively similar across the three AI-involved conditions. This is analytically important because it illustrates a plausible situation in which AI can produce linguistically coherent stories without producing equivalent perceptions of authenticity. The conceptual distinction prevents brands from equating fluent writing with credible storytelling. A narrative can be polished, logically ordered, and emotionally expressive while still provoking doubt about whether it reflects genuine organizational experience.

5.2 Authenticity as a Central Predictor of Skepticism

Across all 160 simulated observations, Perceived Narrative Authenticity and Consumer Skepticism produced a Pearson correlation of approximately $r = -0.732$. Within the constructed model, higher authenticity is therefore associated strongly with lower skepticism.


This relationship is theoretically consistent with brand-authenticity scholarship, which identifies credibility, integrity, continuity, and symbolic meaning as central dimensions of authentic brand perception. It is also consistent with broader consumer research showing that authenticity judgments concern whether marketplace offerings are interpreted as genuine and capable of supporting valued consumer identities. However, the numerical coefficient itself is generated by simulation and should not be interpreted as an empirical estimate.

The result implies that future research should avoid treating AI disclosure as the only explanatory variable. A consumer may know that a narrative was AI-generated yet still perceive it as authentic when the underlying facts are verifiable, the story corresponds to brand history, and humans remain responsible for the final communication. Conversely, a human-authored narrative can still produce skepticism when it exaggerates purpose, misrepresents customers, or manufactures emotional sincerity.

5.3 Graphical Interpretation of Authenticity and Skepticism

The scatter plot below illustrates the relationship between authenticity and skepticism across all four simulated source conditions. Each point represents one synthetic consumer evaluation, while the fitted trend line summarizes the overall direction of association.

Graph 1: Perceived Authenticity and Skepticism Across Brand Narrative Sources



The graph shows substantial overlap between conditions. Some disclosed AI-assisted narratives receive authenticity scores comparable with human-authored narratives, and some human-authored narratives produce elevated skepticism. This overlap is conceptually important because it argues against deterministic claims that AI authorship alone makes a narrative inauthentic.

The most skeptical observations are concentrated among cases with low authenticity, particularly in the AI-generated/no-human-review condition. The fitted relationship therefore supports the analytical proposition that perceived authenticity may mediate consumer responses to AI involvement. Future experiments should test this mediation directly.

6. Discussion

6.1 AI Authorship Does Not Automatically Eliminate Narrative Credibility

The first important implication is that artificial-intelligence authorship should not automatically be equated with consumer rejection. The strongest pre-2021 adjacent evidence comes from automated journalism, where Graefe and Bohlken's meta-analysis found no overall credibility disadvantage for automated news relative to human-written news, although human writing performed better on quality and especially readability. Clerwall's earlier experiment likewise found only limited significant differences between automated and journalist-written factual content. These findings caution against simplistic assumptions that consumers always distrust machine-generated language.

However, brand narratives differ from routine data-based reporting in the type of meaning they communicate. A financial earnings summary or sports result can be judged largely through factual accuracy, whereas a founder story, purpose statement, customer experience, or community narrative often makes claims about intention, emotion, identity, and values. Algorithm-aversion research indicates that reactions to algorithms depend on task characteristics, supporting the hypothesis that consumers may react differently when AI performs subjective or socially interpretive tasks.

The strategic implication is that brands should identify which components of storytelling can be safely automated. AI may be suitable for ideation, structure development, language adaptation, summarization, variation testing, or drafting from verified source material. Claims involving lived experience, personal testimony, historical events, ethical commitments, social identity, or emotionally sensitive topics warrant stronger human involvement because authenticity depends on the relationship between words and actual experience.

6.2 Disclosure and Human Review as Accountability Signals

The second implication concerns transparency. Graefe and Bohlken found that automated-content labeling can influence audience assessments even when the content itself remains constant. This creates a difficult managerial temptation: if disclosure sometimes lowers ratings, brands may prefer not to disclose AI involvement. Yet maximizing short-term message effectiveness through concealment can create a longer-term trust problem if consumers later discover that a narrative presented as personally authored or experientially grounded was substantially generated by an automated system.

Disclosure should therefore be proportionate to the significance of AI's contribution. Minor grammar correction does not carry the same authorship significance as autonomous generation of a founder letter, customer testimonial, employee story, or brand-purpose narrative. A responsible policy can distinguish supportive uses from material generative authorship rather than applying either complete silence or excessive labeling to every automated function.

Human review is equally important because disclosure without accountability is insufficient. The simulation assigns lower skepticism to AI-generated narratives with human review than to fully unsupervised AI generation. This difference represents a theoretical proposition: consumers may be more accepting when the brand clearly assumes responsibility for accuracy, tone, and meaning. In practical governance terms, the organization—not the AI system—must remain accountable for published communication.

6.3 Authenticity Should Be Governed at the Evidence Level

The third implication is that responsible AI storytelling requires evidence governance, not merely stylistic review. Brand authenticity research conceptualizes credibility and integrity as central components of authentic brand perception. A brand therefore cannot preserve authenticity simply by instructing an AI system to "sound authentic." Authenticity depends on whether the narrative is grounded in truthful source material and consistent with actual organizational behavior.

This becomes especially important for quotations, testimonials, employee stories, founder experiences, community relationships, sustainability narratives, and social-impact communication. AI systems can generate plausible experiential details, but plausible language should never be presented as factual human testimony unless the underlying experience has been verified. Synthetic dialogue should be identified as illustrative where used. Similarly, an AI-generated brand-history narrative should be checked against reliable archival or organizational records rather than evaluated only for stylistic consistency.

The same principle applies to emotional tone. Narrative-transportation research demonstrates that stories can influence consumers by drawing them cognitively and emotionally into a narrative world. Because this persuasive mechanism can be powerful, brands have a heightened responsibility not to use AI-generated emotional detail to simulate experiences that never occurred. Responsible storytelling should therefore preserve a clear connection between narrative persuasion and evidentiary truth.

7. Conclusion

Artificial intelligence creates substantial opportunities for brand communication by enabling rapid drafting, variation, personalization, and linguistic adaptation. Yet brand narratives should not be evaluated as ordinary content-production outputs. Their persuasive value frequently depends on perceived authenticity, credibility, continuity, symbolic meaning, and emotional connection. Pre-2021 research on narrative transportation and brand authenticity provides a strong theoretical basis for expecting these qualities to influence consumer responses. At the same time, automated-journalism research demonstrates that machine-produced text does not necessarily suffer an inherent credibility deficit, making a deterministic "AI equals distrust" interpretation theoretically weak.

The simulation developed in this study provides a transparent analytical illustration of how different AI-authorship structures might influence consumer skepticism. Mean skepticism increased from 36.2 for Human-Authored narratives to 43.7 for disclosed AI-Assisted narratives, 53.6 for AI-Generated/Human-Reviewed narratives, and 63.0 for AI-Generated narratives without human review. Perceived Narrative Authenticity and Consumer Skepticism generated a modeled correlation of $r =$

–0.732. These values are synthetic and must not be interpreted as empirical estimates of actual consumer behavior.

The principal theoretical contribution is the proposition that consumer skepticism toward AI-generated brand narratives should be understood as an authenticity-and-accountability problem rather than a technology-rejection problem. AI involvement may become acceptable when the resulting narrative remains faithful to verifiable brand identity, factual claims are checked, material automation is disclosed appropriately, and identifiable humans assume responsibility for the final communication. Skepticism should therefore not be treated merely as friction that marketing technology must overcome. In many situations, skepticism performs a useful consumer-protection function by encouraging audiences to question whether apparently intimate or emotionally persuasive brand communication reflects genuine experience.

Future empirical research should test the framework using randomized experiments that hold narrative text constant while varying authorship disclosure and human-review information. Researchers should compare utilitarian and emotional brand stories, established and unfamiliar brands, factual and purpose-oriented narratives, low- and high-involvement products, and consumers with different levels of AI familiarity. Studies should additionally measure general advertising skepticism, prior brand trust, perceived social presence, narrative transportation, authenticity, disclosure fairness, purchase intention, and willingness to share the narrative. Longitudinal research would be particularly valuable for determining whether repeated exposure to AI-generated storytelling reduces skepticism through familiarity or increases it when consumers perceive declining human accountability. The strongest responsible-marketing position is therefore not to conceal the use of artificial intelligence, nor to reject it categorically, but to use it within a communication system where technological efficiency remains subordinate to truthfulness, traceable evidence, human responsibility, and authentic brand meaning.

References

1. Aaker, J. L. (1997). Dimensions of brand personality. *Journal of Marketing Research*, 34(3), 347–356.
2. Aaker, J. L., Fournier, S., & Brasel, S. A. (2004). When good brands do bad. *Journal of Consumer Research*, 31(1), 1–16.
3. Campbell, M. C., & Kirmani, A. (2000). Consumers' use of persuasion knowledge: A review and directions for future research. *Journal of Consumer Research*, 27(1), 1–31.
4. Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, 56(5), 809–825.
5. 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, 24–42.
6. Deighton, J., & Kornfeld, L. (2017). Digital marketing, data and consumer privacy. *Harvard Business Review*, 95(6), 1–8.
7. Flavián, C., Ibáñez-Sánchez, S., & Orús, C. (2019). The impact of virtual, augmented and mixed reality technologies on the customer experience. *Journal of Business Research*, 100, 547–560.



8. Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research*, 21(1), 1–31.
9. Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172.
10. Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49, 30–50.
11. Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, 46(4), 629–650.
12. Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing Science*, 38(6), 937–947.
13. Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing*, 85(1), 131–151.
14. Sundar, S. S., & Kim, J. (2019). Machine heuristic: When we trust computers more than humans. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13.
15. Wang, W., & Benbasat, I. (2005). Trust in and adoption of online recommendation agents. *Journal of the Association for Information Systems*, 6(3), 72–101.
16. Wood, S., & Moreau, C. P. (2006). From fear to loathing? How emotion influences the evaluation and early use of innovations. *Journal of Marketing*, 70(3), 44–57.
17. Zhang, J., & Mao, E. (2016). From online motivations to ad clicks: A sequential explanation of the effects of motivations to use social media on advertising responses. *Computers in Human Behavior*, 59, 312–321.
18. Batra, R., & Keller, K. L. (2016). Integrating marketing communications: New findings, new lessons, and new ideas. *Journal of Marketing*, 80(6), 122–145.
19. Grewal, D., Roggeveen, A. L., & Nordfält, J. (2017). The future of retailing. *Journal of Retailing*, 93(1), 1–6.
20. Campbell, M. C., & Kirmani, A. (2008). I know what you're doing and why you're doing it: The role of consumer inference in persuasion. *Journal of Consumer Psychology*, 18(4), 296–304.