ORIGINAL RESEARCH **Open access**

Brand Credibility under Synthetic Content Abundance When Consumers Cannot Reliably Distinguish Organizational Commitment from Automated Persuasion

Sofia Marchetti[1*], Mei Lin[2], Aisha Bello[3], Hannah Brooks[4]

Abstract

Generative artificial intelligence enables organizations to produce persuasive content at a scale and sophistication that can separate observable communication quality from the effort, sacrifice, and commitment once associated with it. This article develops a mechanism-based explanation of that separation. It argues that synthetic content abundance can compress the diagnostic value of persuasive cues when those cues become replicable without equivalent organizational commitment. The resulting credibility effect is conditional: AI provenance can trigger aversion, appreciation, indifference, or positive response depending on task fit, disclosure, prior trust, perceived motives, and accompanying organizational evidence. Credibility dilution is therefore conceptualized as declining cue–commitment diagnosticity rather than a direct penalty for AI use. Verifiable actions, accountable commitments, and word–deed consistency can preserve or restore diagnosticity. This distinction identifies when synthetic persuasion should alter brand credibility and when observed effects are better explained by source preferences, privacy concerns, task perceptions, or other rival mechanisms.

Keywords Generative artificial intelligence, Brand credibility, Synthetic content, Costly signaling, Organizational commitment, Automated persuasion

*Correspondence:
Sofia Marchetti
sofia.marchetti@unibocconi.it
[1] Department of Marketing, Bocconi University, Milan, Italy
[2] Department of Marketing, National University of Singapore, Singapore, Singapore
[3] Department of Business Administration, University of Lagos, Lagos, Nigeria
[4] Department of Marketing, Warwick Business School, Coventry, United Kingdom

Introduction

Brand communication relies on observable cues from which consumers infer less observable qualities of an organization. Generative artificial intelligence unsettles this relationship because persuasive artifacts can appear coherent, tailored, emotionally resonant, and professionally produced without requiring the same organizational effort historically associated with those qualities. Research on digital authenticity identifies growing difficulty in determining whether apparently authentic digital outputs correspond to the actor, process, or organizational reality they imply [1]. Experimental evidence in charitable advertising similarly shows that consumers may not reliably recognize AI-generated imagery and that awareness of synthetic provenance can alter responses [2]. The immediate problem is therefore inferential: an artifact may reveal much about its communicative quality while revealing less about the organization behind it.

Generative systems expand the capacity to produce, vary, analyze, and distribute marketing content at low marginal effort [3]. Artificial intelligence can also assume or augment marketing tasks previously performed through human judgment and

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

organizational labor [4]. Yet consumer effects are heterogeneous. A systematic review of AI-enabled engagement finds substantial variation across technologies, contexts, and psychological processes rather than a stable positive or negative response to AI as a category [5]. A theory of synthetic-content credibility must consequently explain why apparently similar automated communications sometimes damage authenticity or trust, sometimes leave them unchanged, and sometimes remain effective.

This article addresses that problem through mechanism-based theory. Its question is not whether consumers like AI or whether automation improves productivity. It is whether visible persuasive cues continue to diagnose an organization's underlying commitment when comparable cues can be generated synthetically at scale. Organizational commitment is used here in an inferential sense: the willingness of the organization to incur effort, accept accountability, sustain word–deed consistency, or bind itself through actions costly to imitate without the underlying intent.

This focus distinguishes the article from debates on algorithm aversion, chatbot acceptance, virtual-influencer effectiveness, and generic trust in AI. Those literatures identify rival mechanisms and boundary conditions, but they do not themselves explain whether synthetic replication weakens the relationship between visible persuasion and hidden organizational commitment. The proposed account is therefore conditional: credibility should dilute only when automation makes a previously diagnostic cue easier to reproduce without equivalent commitment and no compensating evidence restores its diagnostic value.

Brand Credibility and Costly Signals

Consumers rarely observe organizational motives, capabilities, or future behavior directly. They infer them from communication, prior conduct, reputation, and signals whose credibility depends partly on how difficult they are to imitate deceptively. Costly-signaling research shows that communication can appear more honest when associated with costs, verifiability, or self-sacrificial consequences that are less attractive to deceptive senders [6]. Cost does not guarantee honesty; it can instead increase diagnosticity by changing the feasibility of producing the same signal without the claimed underlying quality.

Brand credibility has relational consequences. It is associated with positive word of mouth through mechanisms connecting brands to consumers' social and self-related evaluations [7]. Authenticity research likewise shows that consumers respond to perceived genuineness and consistency, with effects varying by self-congruence [8]. Credibility, authenticity, and identification are therefore related but nonidentical. Credibility concerns whether claims and implied qualities are believable; authenticity concerns perceived genuineness and consistency; identification concerns the consumer's relation to the brand.

Consumers also infer what communication reveals about organizational motives and conduct. Experimental work shows that transparency can strengthen inferred brand integrity through empowerment and perceived consistency when the information is meaningful rather than merely symbolic [9]. Influencer research similarly finds that audiences distinguish intrinsic passion and transparent commerciality from self-presentation attributed mainly to external reward [10]. Visible communication is thus evaluated partly as evidence about what stands behind it.

This logic separates persuasive performance from commitment-bearing evidence. A message can be eloquent, personalized, emotionally appropriate, and aesthetically convincing while revealing little about whether the organization invested resources, accepted risk, changed behavior, or undertook obligations consistent with what it implies. If synthetic systems reduce the effort required to reproduce the persuasive surface of a signal, the theoretical issue becomes whether they also reduce the consumer's ability to infer the organizational state that once made the signal credible.

Synthetic Content Abundance

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

Recent experimental studies show that consumer responses to synthetic brand content vary with authenticity judgments, inferred motives, and message or role fit [11–13]. GenAI-produced social content can reduce perceived brand authenticity, yet the effect is attenuated when AI assists rather than replaces human creation [11]. AI-generated diversity cues can activate identity-threat processes when organizational motives appear inauthentic [12], whereas AI-ad effectiveness can improve or deteriorate depending on appeal type, consumer efficacy, and the social role assigned to AI [13]. Synthetic provenance therefore does not carry one psychological meaning.

The theoretical importance of abundance lies in replication capacity. As generative tools produce many variants of persuasive content, the scarcity of polished expression, personalized tone, visual sophistication, and rapid response declines. A cue that once implied substantial human effort may become obtainable without the same organizational investment. This does not make the cue false; it changes what the cue warrants inferring.

AI advertising also raises questions about who or what is speaking, persuading, and bearing responsibility [14]. Under synthetic abundance, provenance alone is insufficient. A firm may use AI to express a deeply held commitment, efficiently communicate an already enacted practice, or simulate concern at negligible organizational cost. Synthetic content abundance is therefore defined here as an environment in which persuasive cues can be reproduced, varied, and distributed with sufficiently low production friction that their presence becomes less informative about the organizational resources, intentions, and obligations behind them. This is a theoretical environmental condition, not a claimed market-wide effect size.

Automated Persuasion and Signal Replication

Automated persuasion shows that persuasive functionality can be reproduced without preserving the source properties historically associated with it. AI influencers can generate meaningful consumer responses while altering judgments concerning agency, trust, and psychological distance [15]. Other experiments show that AI endorsers can transfer evaluations to brands and sometimes perform comparably with human endorsers [16]. Persuasive effectiveness therefore need not imply identical beliefs about human effort, intention, or accountability.

Virtual-influencer studies further show that source effects vary with origin disclosure, emotional narrative, product–endorser fit, and organizational affiliation [17–19]. The persuasive agent has no fixed meaning outside these relationships. Consumers interpret it through fit, narrative, affiliation, and the relationship between the synthetic representative and the brand. Privacy creates a distinct route: virtual-influencer ecosystems can generate data-vulnerability concerns even when no inference about organizational commitment is involved [20].

Conversational agents provide a second replication setting. Chatbot responses vary with regulatory focus, privacy concerns, social cues, task complexity, service criticality, and disclosure conditions [21–23]. Field evidence also shows that early chatbot disclosure can sharply reduce purchasing even when the system performs effectively [24]. These findings prevent attribution of every adverse response to weakened organizational commitment. Machine-category stereotypes, perceived empathy, privacy, and task mismatch can produce similar observed outcomes.

For this theory, automated persuasion matters when it creates surface functional equivalence while preserving underlying nonequivalence. An AI system can reproduce a tone of care, inclusion, expertise, enthusiasm, or responsiveness, so the cue may still perform its immediate communicative function. What potentially changes is its production cost and evidentiary value. The proposed mechanism begins at this separation between what communication accomplishes and what its production warrants inferring about the organization.

Organizational Commitment as an Unobservable Attribute

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

The hidden attribute in this theory is not sincerity as an inner psychological state or authenticity as a global consumer judgment. Organizational commitment denotes the extent to which the firm has undertaken effort, resource allocation, behavioral change, accountability, or obligation consistent with its communication. Consumers usually infer this commitment from traces such as costly actions, sustained consistency, transparent practices, credible sacrifices, or arrangements difficult to maintain if the professed commitment were merely rhetorical.

AI provenance cannot identify that state. An experiment on AI-disclosed brand voice found no necessary disadvantage for AI-generated communication in perceived brand authenticity or brand attitude [25]. This null result rules out a deterministic AI-to-credibility path. Automated communication can remain credible when it aligns with other evidence about the organization or when consumers do not treat automation as relevant to the focal judgment.

Prior organizational trust also conditions AI acceptance. Evidence across multiple companies and service industries indicates that trust in the company supports AI-service adoption while AI autonomy changes that relationship [26]. Research on voice-controlled AI likewise finds that trust, risk, novelty, and brand involvement jointly shape evaluations [27]. Consumers therefore approach synthetic communication with priors about the firm, technology, task, and consequences of delegation.

The mechanism developed here treats commitment as latent and cue–commitment diagnosticity as variable. Diagnosticity is high when observing a cue materially increases the reason to believe that the organization has undertaken the commitment implied by the message. It falls when the same cue can be reproduced with little corresponding effort or when its production becomes too ambiguous to distinguish enacted commitment from persuasive simulation. Low diagnosticity does not mean the message is false; it means the message carries less evidentiary weight about the organization behind it.

Credibility can consequently be preserved when automated persuasion is accompanied by difficult-to-fake organizational actions, accountable promises, verifiable disclosures, or sustained word–deed consistency. The managerial question is not whether to eliminate AI-generated persuasion, but whether consumers retain credible evidence connecting what is said to what the organization has committed to do.

The same persuasive output can therefore imply different levels of organizational commitment depending on what produced it, what accompanies it, and what consumers can verify. **Table 1** distinguishes replicated persuasive functions from the commitment-bearing evidence needed to interpret them without assuming that automation itself is either authentic or inauthentic.

Table 1. When replicated persuasion ceases to diagnose organizational commitment

Replicated persuasive function	Observable surface cue	Why diagnosticity can fall	Commitment-bearing evidence that can restore diagnostic value	Rival mechanism kept separate	Status
Care and empathy	Warm, responsive, emotionally appropriate language	Empathetic phrasing can be generated without corresponding service effort or organizational sacrifice	Service-recovery authority, costly remedies, human escalation, sustained response behavior	Machine-interaction preference; perceived humanness	Proposed synthesis
Inclusion and social concern	Diverse representation and prosocial claims	Symbolic representation can be reproduced	Verifiable practices, enduring resource	Identity threat; tokenism attribution	Proposed synthesis

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

		independently of resource allocation or internal practice	commitments, accountable goals		
Expertise and competence	Fluent recommendations and confident technical language	Generative fluency can reproduce expertise cues without demonstrating accountable knowledge or decision quality	Traceable expertise, decision ownership, validation and correction mechanisms	Task-dependent algorithm appreciation or aversion	Proposed synthesis
Brand voice and relational consistency	Stable tone and personalized wording	Stylistic continuity can be automated without sustained human attention or organizational continuity	Longitudinal word–deed consistency, stable policies, accountable claims	Neutral or positive response to AI-disclosed brand voice	Proposed synthesis
Responsiveness and personalization	Fast replies, tailored offers, individualized recommendations	Rapid adaptation can be scaled without additional care or willingness to act on consumer needs	Meaningful service discretion, actual problem resolution, privacy-respecting data practice and recourse	Privacy concern, regulatory focus, task complexity	Proposed synthesis

Note: The table separates observable persuasive performance from evidence bearing on the latent organizational commitment inferred from that performance. “Proposed synthesis” denotes relationships integrated by this article rather than a complete causal chain directly validated in any single study.

Credibility Dilution Mechanisms

The proposed mechanism begins when synthetic systems alter the cost structure of producing persuasive cues. A polished response, empathetic tone, visually inclusive campaign, or consistent brand voice may once have required visible human labor, specialist expertise, or sustained organizational attention. Generative systems can now reproduce some of those outputs with far less incremental effort. Cost reduction alone does not reduce credibility. It becomes theoretically relevant when consumers had previously used the costly production of the cue as evidence that the organization possessed or had undertaken the quality the cue appeared to express.

This mechanism must be separated from generic machine preference. Research on algorithm appreciation shows that people sometimes rely more on algorithmic than human advice [28]. Conversely, algorithm aversion increases for tasks perceived as subjective and can change with task framing and human-likeness [29]. These findings imply that an unfavorable response to AI cannot, by itself, establish credibility dilution. A source-category response predicts evaluation from the fact that a machine is involved. The cue–commitment mechanism predicts a different pattern: credibility should change when the same observable cue becomes easier to produce without the corresponding organizational commitment, even when AI provenance is held constant.

A second process is provenance ambiguity. Synthetic abundance makes it harder to infer who authored a message, what organizational resources were committed to it, and whether persuasive performance reflects durable organizational capability or an easily replicated output. Consumer experiences with AI are already known to differ across data capture, classification, delegation, and social interaction [30]. Strategic marketing theory similarly distinguishes mechanical,

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

thinking, and feeling AI because task requirements and relational expectations differ [31]. Cue replication should therefore matter most where consumers normally read communication as evidence of relational effort, moral commitment, specialized judgment, or accountable attention.

A third process is authenticity attenuation. When consumers infer that a seemingly expressive cue required little organizational involvement, the cue may carry less evidence about genuine organizational intention. The observed penalty in GenAI social content is consistent with this possibility [11], but the null effect for AI-disclosed brand voice demonstrates that it is not sufficient [25]. Authenticity attenuation should occur only when automation removes something consumers regard as diagnostic and no alternative evidence compensates for it.

A fourth process is misattribution. Disclosure can reduce trust in some service settings [23], and early disclosure can sharply reduce purchasing in a field environment [24]. Those outcomes may reflect stereotypes about machines, perceived lack of empathy, or task expectations rather than a weakened inference of organizational commitment. Mechanism-based theory therefore requires competing explanations to remain observable. Credibility dilution is supported only when changes in perceived cue diagnosticity mediate the effect after relevant source, task, privacy, and fit explanations have been considered.

The resulting causal claim is deliberately bounded: synthetic persuasion dilutes brand credibility when replication lowers the evidentiary value of cues that consumers use to infer organizational commitment. The effect should weaken when commitment was never inferred from the cue, when consumers value automation for the task, or when verifiable organizational behavior supplies stronger evidence than the synthetic cue itself.

Synthetic-Abundance Credibility Framework

The framework integrates these processes around cue–commitment diagnosticity. Its starting point is not AI use but a persuasive cue that consumers may treat as evidence about the organization. Synthetic abundance places replication pressure on that cue. When replication preserves surface performance while reducing the effort, scarcity, or organizational involvement required to produce it, diagnosticity may decline. If the cue was never interpreted as evidence of commitment, replication should have little credibility consequence.

The central mediator is the consumer's judgment that the observed cue meaningfully distinguishes a committed organization from one capable of producing the same appearance without the corresponding action. Costly-signaling research provides the general logic that verifiability and differential cost can support honesty judgments [6]. Transparency research indicates that integrity inferences strengthen when information provides meaningful evidence and perceived consistency [9]. The framework therefore predicts that synthetic replication is most consequential where it compresses this distinguishing function.

Countervailing mechanisms can interrupt the cascade. A brand may connect automated communication to verifiable practices, accountable service remedies, resource commitments, or behavior that is difficult to reproduce symbolically. Human assistance can also preserve authenticity more effectively than perceived wholesale replacement in some settings [11]. Prior company trust may supply an additional interpretive resource [26], although relying on reputation alone can create vulnerability if later conduct contradicts the message. These mechanisms do not make synthetic content inherently credible; they provide independent evidence that can restore the connection between communication and organizational action.

The framework also allows bifurcation. AI-generated advertising can perform well under fitted appeals and social roles [13]. Virtual influencers can be effective when product–endorser fit or brand affiliation aligns with the interaction [18]. Brand affiliation can also change the relative effectiveness of human and virtual influencers [19]. Chatbot disclosure effects vary with service criticality and failure context [23]. Identical synthetic provenance can therefore lead to different credibility outcomes because provenance is filtered through task, relational, organizational, and service conditions.

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

This integrated account yields three bounded outcome regions. In a preserved-diagnosticity region, automation changes production but commitment remains inferable from stronger accompanying evidence. In a diluted-diagnosticity region, replicated cues remain persuasive yet provide weaker grounds for believing that the organization has made the commitment they imply. In a rival-mechanism region, observed responses are better explained by machine preference, privacy, task fit, identity threat, or other source effects. Movement between regions is possible as organizations add or remove verifiable commitment-bearing evidence.

The proposed cascade therefore depends on a loss of diagnostic connection rather than on automation itself, with countervailing evidence capable of diverting the process before credibility is diluted.

Figure 1 maps the conditional pathways linking synthetic replication, cue–commitment diagnosticity, and credibility.

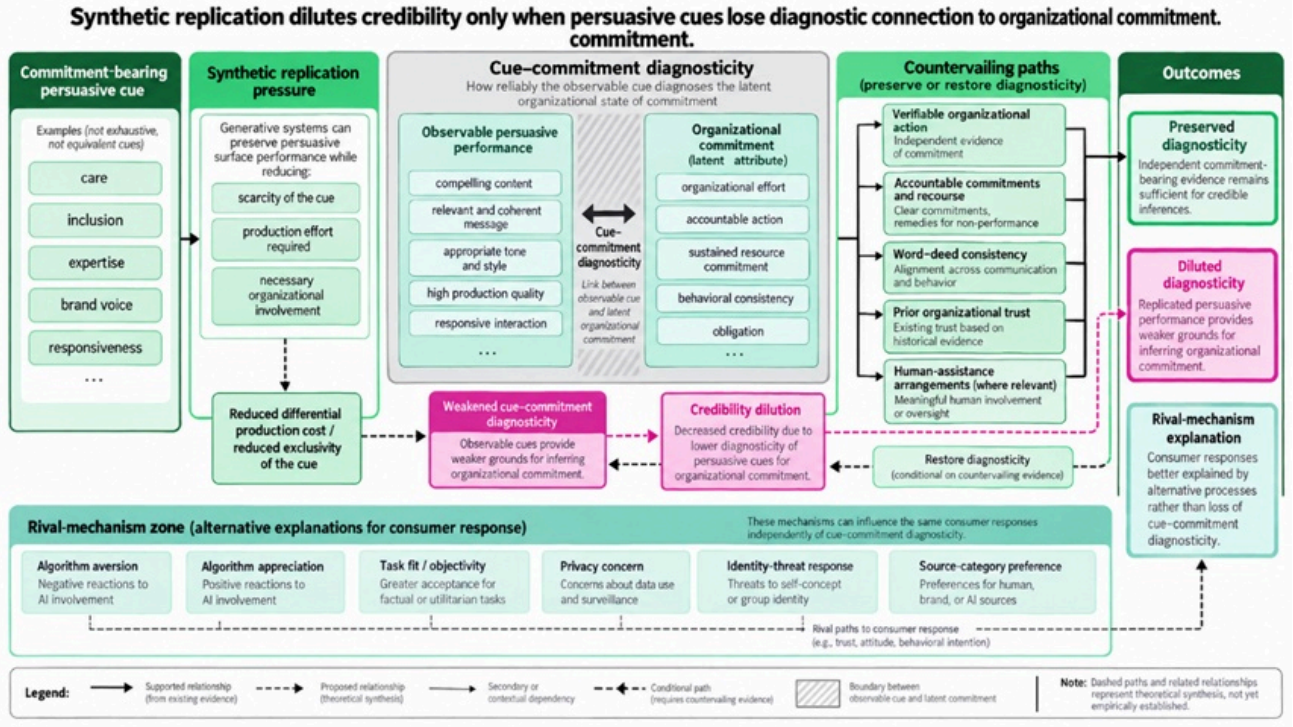

Figure 1. Synthetic replication dilutes credibility only when persuasive cues lose diagnostic connection to organizational commitment

Marketing Propositions

The framework supports propositions that discriminate cue–commitment diagnosticity from simpler AI-source effects.

Proposition 1. Holding AI provenance and content quality constant, brand credibility will be lower when a persuasive cue can be reproduced without the organizational effort, sacrifice, or accountable action that consumers normally treat as evidence of commitment. The proposition follows from costly-signaling logic [6] but concerns the replication of brand cues rather than signal cost in the abstract.

Proposition 2. The adverse credibility effect of synthetic replication will weaken when independent, verifiable organizational evidence preserves cue–commitment diagnosticity. Transparency can support integrity inferences when it is informative and empowering [9], while the absence of an authenticity penalty for AI-disclosed brand voice shows that provenance alone need not reduce credibility [25].

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

Proposition 3. If an observed negative response is primarily driven by algorithm aversion, task subjectivity should explain more variance than commitment-bearing evidence. If cue–commitment diagnosticity is the active mechanism, manipulating verifiable organizational commitment while holding machine provenance constant should shift credibility even within the same task. Task-dependent algorithm aversion provides the principal rival prediction [29].

Proposition 4. Disclosure will improve credibility only when it changes what consumers can infer about organizational conduct, responsibility, or commitment; disclosure that merely labels automation may have neutral or negative effects. Context-dependent disclosure findings in service settings support treating this distinction as conditional [23].

These propositions specify tests capable of weakening the theory. Failure of diagnosticity manipulations to alter credibility, once source category and task perceptions are controlled, would favor rival explanations over the proposed mechanism.

The propositions become discriminating only when each predicted pattern is paired with an observation that would instead favor a rival explanation or place the case outside the proposed mechanism.

Table 2 pairs the proposed credibility mechanisms with discriminating tests and rival explanations.

Table 2. Discriminating tests for synthetic-abundance credibility mechanisms

Theoretical test	Manipulated or contrasted condition	Pattern predicted by cue–commitment diagnosticity	Pattern favoring a rival explanation	Interpretation boundary
Replication without commitment	Same AI provenance and content quality; vary whether the cue requires verifiable organizational effort or action	Credibility changes with commitment-bearing evidence even though provenance is constant	Credibility follows AI-versus-human source categorization regardless of commitment evidence	Distinguishes diagnosticity loss from source-category preference
Restorative evidence	Hold synthetic cue constant; add or remove independently verifiable organizational conduct	Verifiable conduct restores credibility by reconnecting cue and commitment	Added conduct evidence has little effect once another explanatory process is accounted for	Tests the countervailing mechanism rather than disclosure alone
Task dependence	Hold commitment evidence constant across more objective and more subjective tasks	Diagnosticity effect persists where the cue is used to infer organizational commitment	Response is explained mainly by perceived task objectivity or human-versus-algorithm preference	Separates commitment diagnosticity from task-dependent algorithm aversion
Disclosure content	Compare automation labels with disclosures that also reveal responsibility, conduct, or recourse	Disclosure helps when it increases information about organizational commitment	Automation labeling produces the response regardless of commitment-relevant information	Distinguishes commitment-relevant transparency from automation labeling
Temporal consistency	Compare isolated persuasive output with repeated word–deed consistency over time	Accumulating consistent conduct moves cases toward preserved diagnosticity	Credibility remains dominated by source category despite accumulating conduct evidence	Tests whether organizational behavior can restore diagnostic value

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

Note: The table states theory-derived discriminating tests rather than established effect sizes. “Rival explanation” denotes an alternative process capable of producing a similar observed response without reduced cue–commitment diagnosticity.

Managerial Implications

The framework shifts managerial attention from whether AI should be visible to what evidence remains available after automation. Organizations using generative systems for relational, prosocial, expert, or service communication should identify which cues consumers may interpret as evidence of organizational effort or intent. If those cues become cheap to reproduce, managers can preserve diagnosticity by linking them to verifiable conduct: service guarantees with real recourse, transparent policies, accountable claims, documented implementation, or sustained resource commitments.

Disclosure policy should therefore be designed around informational function. Transparency can strengthen integrity when it empowers consumers to assess conduct [9], whereas disclosure can reduce trust or purchasing when it primarily activates unfavorable machine-source expectations [23]. Field evidence likewise demonstrates that disclosure timing can materially change purchasing responses [24]. The managerial objective is not maximal disclosure in every interaction; it is to ensure that consumers can distinguish automated expression from the organizational actions and responsibilities standing behind it.

Automation governance should also reflect task and autonomy. Company trust affects adoption of AI services and changes with the degree of AI autonomy [26]. Strategic deployment should therefore preserve clear organizational ownership where consumers expect judgment, care, or recourse. AI can produce the communication while the organization remains visibly accountable for the commitment.

Conclusion

Synthetic content abundance does not make brand communication inherently incredible. It changes the conditions under which communication can serve as evidence. When persuasive cues become easy to reproduce without the organizational effort, sacrifice, or accountability consumers once inferred from them, those cues can lose diagnostic value even while remaining persuasive.

The resulting problem is best understood as conditional credibility dilution through reduced cue–commitment diagnosticity. Automation may damage, preserve, or improve consumer response depending on task fit, prior trust, source interpretation, and the presence of independent commitment-bearing evidence. Human production is not a guarantee of authenticity, just as AI production is not evidence of insincerity. Brand credibility remains defensible when consumers can still distinguish persuasive appearance from organizational commitment and when the latter is supported by conduct that synthetic persuasion alone cannot cheaply replicate.

Acknowledgements

None

Conflict of interest

None

Financial support

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

None

Ethics statement

None

Received: 06 Feb 2023 Revised: 12 Apr 2023 Accepted: 28 Apr 2023
Published online: 18 March 2024

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

References

1. Pedersen CL, Ritter T. Digital authenticity: Towards a research agenda for the AI-driven fifth phase of digitalization in business-to-business marketing. Ind Mark Manag. 2024;123:162-72. https://doi.org/10.1016/j.indmarman.2024.10.005.

2. Arango L, Singaraju SP, Niininen O. Consumer responses to AI-generated charitable giving ads. J Advert. 2023;52(4):486-503. https://doi.org/10.1080/00913367.2023.2183285.

3. Park HE. The double-edged sword of generative artificial intelligence in digitalization: An affordances and constraints perspective. Psychol Mark. 2024;41(11):2924-41. https://doi.org/10.1002/mar.22094.

4. Davenport TH, Guha A, Grewal D, Bressgott T. How artificial intelligence will change the future of marketing. J Acad Mark Sci. 2020;48(1):24-42. https://doi.org/10.1007/s11747-019-00696-0.

5. Hollebeek LD, Menidjel C, Sarstedt M, Jansson J, Urbonavicius S. Engaging consumers through artificially intelligent technologies: Systematic review, conceptual model, and further research. Psychol Mark. 2024;41(4):880-98. https://doi.org/10.1002/mar.21957.

6. Chaudhry SJ, Wald KA. Overcoming listener skepticism: Costly signaling in communication increases perceived honesty. Curr Opin Psychol. 2022;48:101442. https://doi.org/10.1016/j.copsyc.2022.101442.

7. An J, Do DKX, Ngo LV, Quan THM. Turning brand credibility into positive word-of-mouth: Integrating the signaling and social identity perspectives. J Brand Manag. 2019;26(2):157-75. https://doi.org/10.1057/s41262-018-0118-0.

8. Kumar V, Kaushik AK. Engaging customers through brand authenticity perceptions: The moderating role of self-congruence. J Bus Res. 2022;138:26-37. https://doi.org/10.1016/j.jbusres.2021.08.065.

9. Cambier F, Poncin I. Inferring brand integrity from marketing communications: The effects of brand transparency signals in a consumer empowerment context. J Bus Res. 2020;109:260-70. https://doi.org/10.1016/j.jbusres.2019.11.060.

10. Audrezet A, de Kerviler G, Moulard JG. Authenticity under threat: When social media influencers need to go beyond self-presentation. J Bus Res. 2020;117:557-69. https://doi.org/10.1016/j.jbusres.2018.07.008.

11. Brüns JD, Meißner M. Do you create your content yourself? Using generative artificial intelligence for social media content creation diminishes perceived brand authenticity. J Retail Consum Serv. 2024;79:103790. https://doi.org/10.1016/j.jretconser.2024.103790.

12. Sands S, Demsar V, Ferraro C, Campbell C, Cohen J. Inauthentic inclusion: Exploring how intention to use AI-generated diverse models can backfire. Psychol Mark. 2024;41(6):1396-413. https://doi.org/10.1002/mar.21987.

13. Chen Y, Wang H, Rao Hill S, Li B. Consumer attitudes toward AI-generated ads: Appeal types, self-efficacy and AI's social role. J Bus Res. 2024;185:114867. https://doi.org/10.1016/j.jbusres.2024.114867.

14. Coffin J. Asking questions of AI advertising: A maieutic approach. J Advert. 2022;51(5):608-23. https://doi.org/10.1080/00913367.2022.2111728.

15. Sands S, Campbell C, Plangger K, Ferraro C. Unreal influence: Leveraging AI in influencer marketing. Eur J Mark. 2022;56(6):1721-47. https://doi.org/10.1108/EJM-12-2019-0949.

16. Thomas VL, Fowler K. Close encounters of the AI kind: Use of AI influencers as brand endorsers. J Advert. 2021;50(1):11-25. https://doi.org/10.1080/00913367.2020.1810595.

17. Lim RE, Lee SY. "You are a virtual influencer!": Understanding the impact of origin disclosure and emotional narratives on parasocial relationships and virtual influencer credibility. Comput Hum Behav. 2023;148:107897. https://doi.org/10.1016/j.chb.2023.107897.

18. Kim H, Park M. Virtual influencers' attractiveness effect on purchase intention: A moderated mediation model of the product-endorser fit with the brand. Comput Hum Behav. 2023;143:107703. https://doi.org/10.1016/j.chb.2023.107703.

19. Igarashi R, Bhoumik K, Thompson J. Investigating the effectiveness of virtual influencers in prosocial marketing. Psychol Mark. 2024;41(9):2121-35. https://doi.org/10.1002/mar.22031.

20. Liyanaarachchi G, Mifsud M, Viglia G. Virtual influencers and data privacy: Introducing the multi-privacy paradox. J Bus Res. 2024;176:114584. https://doi.org/10.1016/j.jbusres.2024.114584.

21. Kim WJ, Ryoo Y, Lee SY, Lee JA. Chatbot advertising as a double-edged sword: The roles of regulatory focus and privacy concerns. J Advert. 2023;52(4):504-22. https://doi.org/10.1080/00913367.2022.2043795.

22. Cheng X, Bao Y, Zarifis A, Gong W, Mou J. Exploring consumers' response to text-based chatbots in e-commerce: The moderating role of task complexity and chatbot disclosure. Internet Res. 2022;32(2):496-517. https://doi.org/10.1108/INTR-08-2020-0460.

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

23. Mozafari N, Weiger WH, Hammerschmidt M. Trust me, I'm a bot—repercussions of chatbot disclosure in different service frontline settings. J Serv Manag. 2022;33(2):221-45. https://doi.org/10.1108/JOSM-10-2020-0380.

24. Luo X, Tong S, Fang Z, Qu Z. Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. Mark Sci. 2019;38(6):937-47. https://doi.org/10.1287/mksc.2019.1192.

25. Kirkby A, Baumgarth C, Henseler J. To disclose or not disclose, is no longer the question—effect of AI-disclosed brand voice on brand authenticity and attitude. J Prod Brand Manag. 2023;32(7):1108-22. https://doi.org/10.1108/JPBM-02-2022-3864.

26. Frank DA, Jacobsen L, Søndergaard HA, Otterbring T. In companies we trust: Consumer adoption of artificial intelligence services and the role of trust in companies and AI autonomy. Inf Technol People. 2023;36(8):155-73. https://doi.org/10.1108/ITP-09-2022-0721.

27. Hasan R, Shams R, Rahman M. Consumer trust and perceived risk for voice-controlled artificial intelligence: The case of Siri. J Bus Res. 2021;131:591-7. https://doi.org/10.1016/j.jbusres.2020.12.012.

28. Logg JM, Minson JA, Moore DA. Algorithm appreciation: People prefer algorithmic to human judgment. Organ Behav Hum Decis Process. 2019;151:90-103. https://doi.org/10.1016/j.obhdp.2018.12.005.

29. Castelo N, Bos MW, Lehmann DR. Task-dependent algorithm aversion. J Mark Res. 2019;56(5):809-25. https://doi.org/10.1177/0022243719851788.

30. Puntoni S, Reczek RW, Giesler M, Botti S. Consumers and artificial intelligence: An experiential perspective. J Mark. 2021;85(1):131-51. https://doi.org/10.1177/0022242920953847.

31. Huang MH, Rust RT. A strategic framework for artificial intelligence in marketing. J Acad Mark Sci. 2021;49(1):30-50. https://doi.org/10.1007/s11747-020-00749-9.

© 2024 The Author(s). | Published by Institute for Management, Business, and Accounting Studies