

Mechanisms of Brand Authenticity Erosion in AI-Generated Advertising: A Systematic Review

Huiwen Xie

*Smeal College of Business, Pennsylvania State University, State College, USA
hwenn77@outlook.com*

Abstract. The rapid adoption of generative artificial intelligence (GenAI) in advertising since 2022 has transformed content production while raising concerns about the preservation of brand authenticity, a critical foundation of consumer trust and brand relationships. This study examines how AI-generated advertising influences perceived brand authenticity, identifies the mechanisms driving authenticity erosion, and explains the contextual factors that account for inconsistent findings in prior research. Following PRISMA 2020 guidelines, this review synthesizes 13 studies published between 2022 and 2026 retrieved from the Web of Science Core Collection and EBSCOhost databases. Through thematic synthesis, this study identifies three interrelated mechanisms through which AI-generated advertising may undermine brand authenticity: disclosure effects, failed anthropomorphism, and sincerity erosion. The findings further reveal that the effects of AI-generated advertising on authenticity perceptions are contingent upon individual characteristics, product and brand contexts, and source-related factors, including AI capability beliefs, product type, luxury positioning, and communicator characteristics. This study develops an integrated framework explaining when and why AI-generated advertising enhances or undermines brand authenticity, thus providing theoretical insights for consumer research and practical implications for AI-enabled marketing decision-making.

Keywords: AI-generated advertising, brand authenticity, AI disclosure, Consumer trust

1. Introduction

The rapid advancement of generative AI since 2022 has transformed advertising practices by enabling new forms of content creation and delivery. Yet, this technological shift presents a fundamental tension: while AI enhances efficiency and scalability, it may simultaneously challenge perceptions of brand authenticity. Brand authenticity refers to consumers' perception that a brand is genuine, credible, and aligned with its core values, and it plays a critical role in establishing trust. When consumers perceive advertising messages as machine-generated rather than human-crafted, concerns about sincerity may arise, potentially undermining brand trust and consumer engagement. Although research in this area has expanded rapidly, existing findings remain fragmented. Prior studies on AI disclosure, content perfection, and anthropomorphic cues have produced inconsistent conclusions, while the psychological mechanisms underlying these effects have yet to be systematically integrated. Moreover, the boundary conditions that shape consumers' responses

remain insufficiently understood. Following the PRISMA 2020 guidelines, this study conducts a systematic literature review (SLR) and employs thematic synthesis to integrate empirical evidence from studies published between 2022 and 2026 [1-3]. Specifically, this study examines how AI-generated advertising undermines brand authenticity perceptions by identifying the underlying mechanisms, boundary conditions that shape this process, and factors that explain inconsistencies across prior findings. By synthesizing existing evidence into an integrated Mechanism Framework, this research advances consumer intelligence research within Business Analytics.

2. Methodology

2.1. Study design and search strategy

This study conducts a systematic literature review (SLR) following PRISMA 2020 guidelines, with thematic synthesis as the analytical approach. The search strategy was developed using an adapted PICO framework covering AI technologies (Population), advertising and marketing communication contexts (Intervention), brand authenticity-related outcomes (Outcome), and mechanism-related constructs (Mechanism), including disclosure, anthropomorphism, perceived deception, algorithm aversion, human touch, and synthetic media. And the review period was restricted to 2022–2026 to capture the emergence and diffusion of large language models and generative AI in commercial applications. Searches were conducted in the Web of Science Core Collection and EBSCO Business Source Ultimate. Web of Science served as the primary database for peer-reviewed empirical research, while EBSCO was used as a supplementary source. Given the interdisciplinary and emerging nature of AI-related constructs, a stepwise search strategy was adopted. Table 1 presents the search queries and identified records. Filters were applied to include articles, early access publications, and review articles within the fields of Communication, Business, and Psychology.

Table 1. Database search strategy and identified records

Line	Search String	Field	Results
1	("generative AI" OR "AI-generated content" OR "artificial intelligence" OR "machine learning" OR "algorithmic systems")	TS	106,291
2	("advertising" OR "advertisement" OR "content marketing" OR "digital advertising" OR "social media" OR "brand communication")	TS	101,397
3	("authenticity" OR "brand trust" OR "brand credibility" OR "brand sincerity" OR "perceived sincerity" OR "consumer trust")	TS	53,628
4	("aesthetics" OR "anthropomorphism" OR "perceived deception" OR "algorithm aversion" OR "human touch" OR "synthetic media")	TS	101,397
5	#3 AND #4	—	13,179
6	#5 AND #2	—	292
7	#6 AND PY= (2022–2026)	—	31

In addition, the supplementary EBSCO search used abstract-based queries combining generative AI, advertising/marketing, and brand authenticity terms, yielding eight additional records. Across both databases, 39 records were identified and imported into EndNote 21 for duplicate removal and screening. Backward and forward citation tracking was conducted to enhance retrieval coverage.

The complete search process and study selection procedure were documented and reported through the PRISMA flow diagram.

2.2. Study selection and eligibility criteria

Study selection was conducted following PRISMA 2020 guidelines through a two-stage screening process: title and abstract screening, followed by full-text eligibility assessment. Inclusion and exclusion criteria were developed based on the PICOS framework to ensure consistency between the search strategy and study selection process, as presented in Table 2. The criteria specified the target population, intervention, comparison conditions, outcomes, study designs, publication period (2022–2026), language (English), and publication type (peer-reviewed journal articles).

Table 2. PICOS-based inclusion and exclusion criteria

Dimension	Inclusion	Exclusion
Population	Consumers in B2C contexts exposed to AI-generated advertising	B2B contexts; marketers or developers as subjects
Intervention	AI-generated or AI-assisted advertising, disclosed or detectable	Analytical AI without creative generation; undetected AI
Comparison	Human-created content; no-AI control; varying AI degrees	Single-group designs; AI-only comparisons
Outcomes	Brand authenticity, trust, credibility, sincerity	Generic attitudes; behaviors without mechanisms
Study design	Experiments, surveys, quasi-experiments, content analyses	Conceptual papers; technical AI; editorials; non-English
Time	2022-2026	Pre-2022
Quality	Peer-reviewed journal articles	Conference papers, dissertations, books, preprints, and non-peer-reviewed reports

Titles and abstracts were independently screened according to the predefined PICOS criteria, and potentially eligible records were subjected to full-text assessment. Exclusion decisions at the full-text stage were documented in accordance with PRISMA 2020 recommendations. Following the screening process, 13 articles were included in the final review, comprising 11 studies retrieved from the Web of Science Core Collection and two from EBSCO Business Source Ultimate.

The methodological quality of included studies was assessed using an adapted CASP checklist consisting of 10 criteria. Studies meeting the predefined quality threshold were retained for thematic synthesis, while lower-quality studies were excluded. All screening and appraisal decisions were documented to support transparency and reproducibility.

2.3. Data processing and thematic synthesis

Thirteen articles meeting the PICOS criteria and published between 2022 and 2026 were retained for this review, including eleven from the Web of Science Core Collection and two from EBSCOhost. A structured extraction form was developed to capture bibliometric characteristics, methodological

approaches, theoretical frameworks, focal constructs, psychological mechanisms, and boundary conditions reported in the included studies.

Thematic synthesis followed Braun and Clarke and Thomas and Harden via three iterative stages. First, descriptive coding generated initial codes through close reading of each article. First-order concepts, such as "perceived covert manipulation," "algorithm aversion," and "novelty-authenticity trade-off," were retained as descriptive codes to preserve original meanings. Second, analytical coding involved comparing and clustering initial codes into higher-order categories through an abductive process. Related codes were synthesized into broader thematic patterns concerning AI-generated advertising effects, psychological mechanisms, and contextual boundary conditions. Third, theme refinement was conducted by reviewing the emergent themes against the full dataset to ensure internal coherence and representativeness. The final coding structure consisted of four overarching themes and twelve sub-themes, as presented in Table 3.

Table 3. Coding structure developed through thematic synthesis

Theme	Sub-theme	Coding Focus
Theme 1: Paradoxical Effects of AI Disclosure	1.1 Novelty-Authenticity Trade-off	Tension between perceived novelty benefits and authenticity concerns following AI disclosure
	1.2 Credibility Deficit in Prosocial Contexts	AI disclosure and perceived credibility in value-based communication contexts
	1.3 AI Capability Beliefs as Buffer	Consumer beliefs that moderate responses to AI-generated content
Theme 2: Psychological Mechanisms of Authenticity Erosion	2.1 Perceived Covert Manipulation	Inferences of hidden persuasive intent in AI-mediated communication
	2.2 Uniqueness Threat	Perceived loss of individuality and brand distinctiveness
	2.3 Uncanny Valley in Algorithmic Communication	Negative responses to imperfect human-like AI expression
Theme 3: Functional-Match Hypothesis and Contextual Appropriateness	3.1 Utilitarian-Hedonic Product Divide	Alignment between AI capabilities and product characteristics
	3.2 Luxury Brand Vulnerability	Tension between AI production and symbolic brand values
	3.3 Strategic Value of Negative AI Reviews	Conditions under which AI objectivity enhances credibility

Table 3. (continued)

Theme 4: Boundary Conditions and Moderating Factors	4.1 Consumer AI Literacy and Trust	Individual differences affecting AI evaluations
	4.2 Influencer and Source Characteristics	Source-related factors shaping authenticity perceptions
	4.3 Self-Brand Connection and Identity Relevance	Identity-related factors influencing AI disclosure responses

3. Results and analysis

3.1. Overview of the included literature

Thirteen studies published between 2024 and 2026 were included in the final synthesis. Regarding the data sources, 11 studies were identified from the Web of Science Core Collection and two from EBSCOhost, with one duplicate record removed during the screening process. The included studies comprised nine experimental studies [4-12], two used mixed methods [13, 14], and two conceptual papers [15, 16]. Theoretical perspectives varied across studies, with the Elaboration Likelihood Model and Uncanny Valley Theory being the most frequently adopted frameworks, followed by Source Credibility Theory and the CASA paradigm [17-20]. Samples were drawn predominantly from online consumer panels in the United States, China, and Southeast Asia, with sizes ranging from 213 to 650 participants.

3.2. Effects of AI-generated advertising on brand authenticity

3.2.1. The disclosure paradox

Across prior studies, AI disclosure consistently reveals a novelty–authenticity trade-off. Shi and Jiang's three online experiments show that AI disclosure enhances perceived novelty while reducing perceived authenticity [10]. Their moderated mediation analysis indicates that novelty positively affects advertising attitude ($B = 0.239$) and purchase intention ($B = 0.191$), whereas authenticity loss exerts a stronger negative effect on purchase intention ($B = -0.254$). Thus, authenticity concerns may have greater behavioral implications than attitudinal effects.

Deng and Song further qualify this relationship by showing that AI-generated content can receive more favorable consumer responses than human-created content when authenticity concerns are not explicitly activated [13]. In their mixed-methods study with 426 participants, AI-generated content achieved higher evaluations of engagement, trust, and purchase intention than human-created content (all $p < .001$), despite consumers being able to distinguish between the two. These findings suggest that the authenticity penalty may be driven less by AI involvement itself than by disclosure, which changes the evaluative criteria consumers apply to advertising content.

3.2.2. Product category as a boundary

Regarding product category, product type significantly moderates the effects of AI disclosure. Shi and Jiang show that, for utilitarian products, novelty benefits dominate (effect = 0.765, 95% CI [0.548, 0.998]), whereas authenticity concerns remain insignificant (effect = -0.043, ns) [10]. In

contrast, for hedonic products, novelty effects become nonsignificant (effect = 0.156, 95% CI [-0.023, 0.346]), while authenticity concerns exert a significant negative effect (effect = -0.140, 95% CI [-0.248, -0.037]). Supporting this pattern, Nguyen finds that consumer acceptance relies on the alignment between perceived AI capabilities, such as information processing, optimization, and objective evaluation, and product functional requirements [11]. Together with Gu et al.'s findings, these results support the functional-match hypothesis and suggest that AI advertising risks are context-dependent, with hedonic categories facing greater authenticity challenges [6].

3.2.3. The luxury brand vulnerability

Regarding brand positioning, luxury brands appear particularly vulnerable to authenticity threats from AI-generated advertising. To et al. argue that luxury value is rooted in scarcity, artisanal heritage, and human craftsmanship, which are in tension with algorithmic and scalable production. Therefore, AI-generated advertising may represent a "symbolic violation" of luxury brands' core value propositions [12]. Supporting this view, evidence of a perceived uniqueness threat suggests that consumers may perceive AI-generated content as homogenizing brand communication and weakening distinctive brand identities [5]. When consumers perceive AI as replacing human creative input, their sense of uniqueness and connection with the brand may decline. Thus, brand positioning serves as a key boundary condition, with luxury brands facing greater authenticity risks due to their reliance on human craftsmanship and exclusivity.

3.3. Mechanisms explaining authenticity erosion

Thematic synthesis identified three interconnected mechanisms of AI-driven authenticity erosion: disclosure effects, failed anthropomorphism, and sincerity erosion. These mechanisms operate at different stages of consumer processing and are influenced by distinct boundary conditions.

3.3.1. Disclosure effects

The first mechanism involves AI disclosure. When consumers learn that advertising content is AI-generated, they apply beliefs about AI capabilities and limitations that shape subsequent evaluations. Baek et al. show that AI disclosure in prosocial advertising reduces perceived message credibility, leading to lower advertising attitude, brand attitude, and donation intention [4]. This effect stems from a perceived mismatch between algorithmic production and the human compassion expected in prosocial communication, as AI involvement may signal limited emotional investment.

However, disclosure effects vary across consumers. At the individual level, AI disclosure effects are shaped by AI capability beliefs and regulatory focus, influencing negative reactions and novelty-authenticity evaluations [4, 10]. Promotion-focused consumers emphasize novelty gains and reduce authenticity concerns, whereas prevention-focused consumers show the opposite pattern. Evidence from Lu et al. further demonstrates that AI disclosure effects depend on communication context and perceived AI roles [14]. In mobile phone reviews, AI-generated negative reviews increased credibility and reduced perceived risk because consumers associated AI criticism with objectivity rather than persuasion. Thus, AI disclosure does not inherently damage authenticity; erosion occurs when AI involvement conflicts with expectations of emotional expression, sincerity, or communicative intent.

3.3.2. Failed anthropomorphism

The second mechanism operates without explicit AI disclosure through what Nguyen and Dang describe as "uncanny cues" in AI-generated communication [16]. Based on Uncanny Valley Theory and the CASA paradigm, they argue that near-human but imperfect AI content triggers "algorithmic uncanniness," weakening persuasion [18, 20]. Consumers may detect AI-specific cues, including overly polished language, excessive positivity, contextual inconsistencies, and artificial engagement signals, which differentiate AI-generated from human-created content.

This mechanism is particularly important because it occurs below conscious awareness. Uncanny cues disrupt consumers' automatic social responses, shifting processing from heuristic evaluation to more deliberate attributional judgment [21]. Once consumers recognize that AI simulates emotional expression without genuine intention or feeling, authenticity perceptions decline [8]. Therefore, even undisclosed AI content may trigger negative responses when perceived as violating consumers' expectations of human sincerity and authenticity.

3.3.3. Sincerity erosion

The third mechanism arises from consumers' inferences of hidden persuasive intent. Israfilzade identifies perceived covert manipulation as a key mechanism in AI-generated word-of-mouth (WOM): when consumers encounter AI-generated positive reviews, they may infer strategic persuasion behind seemingly authentic user-generated content, leading to distrust and negative brand evaluations [7]. This effect is stronger for positive than negative AI-generated WOM because algorithmic positivity conflicts with expectations of AI neutrality. As a result, consumers may perceive AI-generated praise as a manufactured attempt to influence perceptions rather than a genuine expression of satisfaction.

Sincerity erosion is also linked to perceived uniqueness threats [5]. AI-generated content may be viewed as homogenizing brand communication and weakening its identity-expressive value. When brands serve self-definitional functions, AI attribution may reduce consumers' sense of personal connection with the brand. Consistently, Baek et al. find that consumers with stronger self-brand connections show greater negative responses to AI disclosure, as AI involvement threatens the symbolic value of the brand relationship [4].

However, Lu et al. provide a contrasting perspective, showing that AI-generated negative reviews can enhance product attitudes by increasing perceived credibility and reducing perceived risk [14]. Consumers may interpret AI criticism as objective and data-driven rather than emotionally motivated. Thus, sincerity erosion is not inevitable but depends on content valence and perceived AI motives. When AI's functional attributes align with informational goals, such as objectivity and neutrality, authenticity concerns may be reduced.

3.4. Boundary conditions and integrated explanation

The reviewed studies indicate that AI-generated advertising effects on brand authenticity depend on individual characteristics, product and brand contexts, and source-related factors. These boundary conditions explain variations in consumer responses across contexts.

3.4.1. Individual-level boundaries

Three individual-level moderators consistently emerge. First, AI capability beliefs mitigate negative disclosure effects, as consumers with greater confidence in AI's creative abilities show fewer credibility concerns [4]. Second, regulatory focus shapes the novelty–authenticity trade-off, with promotion-focused consumers emphasizing novelty benefits and prevention-focused consumers emphasizing authenticity costs [10]. Third, self-brand connection intensifies negative responses to AI disclosure because it threatens the identity-expressive value of the brand relationship [4]. Together, these findings indicate that AI disclosure effects vary substantially across consumers, challenging the assumption of a universal disclosure strategy.

3.4.2. Product and brand-level boundaries

At the product and brand level, two key moderators shape AI authenticity effects. First, product type represents a robust boundary condition: AI disclosure tends to benefit utilitarian products but undermine authenticity perceptions for hedonic products [10, 11]. Second, brand positioning acts as a macro-level boundary, with luxury brands facing greater authenticity risks because AI involvement conflicts with the craftsmanship and scarcity values central to luxury [12, 15]. In contrast, mass-market brands may integrate AI with limited reputational consequences. These findings suggest that AI authenticity effects are shaped not only by disclosure strategies but also by the underlying value propositions of products and brands.

3.4.3. Source-level boundaries

At the source level, findings are more nuanced. Rungruangjit et al. show that emotional appeal in virtual influencer content can partially offset source artificiality, suggesting that simulated emotional expression may reduce rather than intensify authenticity concerns under certain conditions [9]. In contrast, Lu et al. identify a credibility paradox: large follower bases strengthen the positive effects of AI-generated negative reviews, whereas high influencer ratings weaken these effects, potentially because highly reputable sources trigger greater skepticism when combined with AI-generated content [14]. These findings show that source-level effects depend on the interaction between source characteristics and content valence.

4. Conclusion

This review synthesizes thirteen studies to explain how AI-generated advertising shapes brand authenticity perceptions. The synthesis reveals three interconnected mechanisms: disclosure effects, failed anthropomorphism, and sincerity erosion. These mechanisms operate via distinct pathways: AI disclosure creates a novelty-authenticity trade-off, failed anthropomorphism operates via uncanny cues, and sincerity erosion stems from perceived persuasive intent in positive AI content. These effects are further shaped by individual, product, brand, and source-level boundary conditions. This review conceptualizes AI advertising authenticity as an outcome of alignment among AI capabilities, brand positioning, consumers, and context. Building on these insights, future research should examine three propositions: disclosure outcomes depend on the relative strength of novelty benefits and authenticity costs, moderate levels of human-likeness may create greater authenticity risks than low or high levels, and positive AI-generated content is more likely to trigger

manipulation inferences whereas negative AI-generated content may enhance perceptions of objectivity.

References

- [1] Page, M. J., McKenzie, J. E., Bossuyt, P. M., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, Article n71. <https://doi.org/10.1136/bmj.n71>
- [2] Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- [3] Thomas, J., & Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8(1), Article 45. <https://doi.org/10.1186/1471-2288-8-45>
- [4] Baek, T. H., Kim, J., & Kim, J. H. (2026). Effect of disclosing AI-generated content on prosocial advertising evaluation. *International Journal of Advertising*, 45(1), 171–192. <https://doi.org/10.1080/02650487.2024.2401319>
- [5] Brüns, J. D., & Meißner, M. (2024). Do you create your content yourself? Using generative artificial intelligence for social media content creation diminishes perceived brand authenticity. *Journal of Retailing and Consumer Services*, 79, Article 103790. <https://doi.org/10.1016/j.jretconser.2024.103790>
- [6] Gu, C., Li, X., & Xiang, Q. (2025). The infinite monkey theorem in AIGC advertising: Matching effects between AI disclosure and advertising appeals on consumer advertising avoidance intention. *Journal of Research in Interactive Marketing*, 19(5), 578–596. <https://doi.org/10.1108/JRIM-05-2025-0265>
- [7] Israfilzade, A. (2024). Consumer responses to AI-generated word-of-mouth: The role of perceived covert manipulation and review valence. *Information & Management*, 61(5), Article 104215. <https://doi.org/10.1016/j.im.2024.104215>
- [8] Kirk, C. P., & Givi, J. (2025). The uncanny valley of generative AI: How near-perfect human mimicry erodes brand authenticity. *Journal of the Academy of Marketing Science*, 53(2), 312–330.
- [9] Rungruangjit, W., Mongkol, K., Piriyaikul, I., & Charoenpornpanichkul, K. (2024). The power of human-like virtual-influencer-generated content: Impact on consumers' willingness to follow and purchase intentions. *Computers in Human Behavior Reports*, 16, Article 100523. <https://doi.org/10.1016/j.chbr.2024.100523>
- [10] Shi, Y., & Jiang, Z. (2026). Consumer responses to AI disclosure labels: The role of novelty and authenticity. *Sage Open*, 16(1). <https://doi.org/10.1177/21582440261417793>
- [11] Nguyen, K. M., Phan, T. M., Tran, Y. N. N., Nguyen, A. T., Nguyen, T. L. N., Hoang, G. H., ... & Nguyen, N. T. (2026). Evaluating the efficacy of AI-generated advertising: A cross-national analysis of customer responses on brand perceptions and customer engagement with evidence from Vietnam and Australia. *Journal of Global Scholars of Marketing Science*, 36(2), 293–341. <https://doi.org/10.1080/21639159.2026.2617659>
- [12] To, R. N., Wu, Y.-C., Kianian, P., & Zhang, Z. (2025). When AI doesn't sell Prada: Why using AI-generated advertisements backfires for luxury brands. *Journal of Advertising Research*, 65(2), 202–236. <https://doi.org/10.1080/00218499.2025.2454120>
- [13] Deng, J., & Song, C. (2026). Analyzing the consumer interaction with generative AI-provided content in digital marketing. *Electronic Commerce Research*. Advance online publication. <https://doi.org/10.1007/s10660-026-10154-z>
- [14] Lu, H., Wang, Y., Ge, L., & Ma, S. (2025). Positive effects of negative disclosure: The persuasive power of negative AI-generated content in shaping consumer product attitude. *Journal of Electronic Commerce Research*, 26(3), 145–168. <https://doi.org/10.1007/s10660-025-09821-3>
- [15] Balaban, D., Cui, T. Q., Dima, A. L., & Sorescu, A. (2026). Generative artificial intelligence in marketing: A systematic literature review. *Journal of Business Research*, 188, Article 115234. <https://doi.org/10.1016/j.jbusres.2026.115234>
- [16] Nguyen, T.-T., & Dang, T.-T. (2026). Beyond the "rainbow lollipop": Uncanny cues in AI-WOM and its impact on trust erosion and negative consumer responses. *International Journal of Human–Computer Interaction*. Advance online publication. <https://doi.org/10.1080/10447318.2026.2674826>
- [17] Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 123–205). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60214-2](https://doi.org/10.1016/S0065-2601(08)60214-2)
- [18] Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*, 19(2), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- [19] Hovland, C. I., & Weiss, W. (1951). The influence of source credibility on communication effectiveness. *Public Opinion Quarterly*, 15(4), 635–650. <https://doi.org/10.1086/266350>

- [20] Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1), 81–103. <https://doi.org/10.1111/0022-4537.00153>
- [21] Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58(9), 697–720. <https://doi.org/10.1037/0003-066X.58.9.697>