Made by a Machine: AI-Disclosure Labels Lower Purchase Intention for Hedonic Products

Kari Nordmann Ola Nordmann

Draft v7, 14 September 2026. Please do not circulate.

Kari Nordmann, BI Norwegian Business School (kari.nordmann@bi.no)
Ola Nordmann, Norwegian School of Economics

Fictional practice material for a BI course. The study, the data and the authors are made up.

Acknowledgements. We thank participants at the BI Marketing Seminar (May 2026) and the EMAC 2026 conference for helpful comments, and Anne Fjell for research assistance.

Abstract

Brands increasingly use generative AI to write product descriptions, and regulators are moving toward requiring that such content be disclosed. We ask how consumers respond to an "AI-written" label. In a preregistered online experiment (N = 476), a disclosure that a product description was written by AI lowered purchase intention for a hedonic product (a scented candle) but not for a utilitarian product (a phone charger). Perceived authenticity explains this effect: AI-written descriptions of hedonic products feel less authentic, and less authentic products are less desirable. The findings extend work on algorithm aversion to marketing communication and suggest that the cost of AI disclosure depends on what is being sold. Managers selling hedonic products should think twice before letting AI write their copy, or should invest in restoring a sense of human involvement.

Keywords: generative AI, disclosure, authenticity, hedonic consumption, algorithm aversion

Introduction

Generative AI now writes a large share of the product descriptions consumers read online. Retail platforms offer one-click AI copywriting, and small sellers in particular have embraced it. At the same time, disclosure is becoming the norm: the EU AI Act requires that certain AI-generated content be labelled as such, and several platforms have introduced "written with AI" badges. Whether such labels help or hurt sellers is an open and practically important question.

Research on algorithm aversion suggests that consumers are often reluctant to rely on algorithms, especially after seeing them err (Dietvorst, Simmons, & Massey, 2015), and that resistance is strongest when consumers feel that their unique circumstances are neglected (Longoni, Bonezzi, & Morewedge, 2019). Yet people also appreciate algorithmic advice in many domains (Logg, Minson, & Moore, 2019), and aversion depends on whether a task is seen as subjective or objective (Castelo, Bos, & Lehmann, 2019). Less is known about how consumers

respond when AI produces not advice but *marketing content*: the words that describe the product itself.

We propose that the answer depends on the product. For hedonic products, which are consumed for pleasure and experience (Dhar & Wertenbroch, 2000), consumers value the sense that a human cared about the product; an AI-written description signals the opposite and lowers perceived authenticity. For utilitarian products, which are bought for their function, who wrote the description should matter little. In our earlier work (Nordmann & Nordmann, 2024), we found that source cues in product descriptions shape quality inferences; here we turn to AI disclosure specifically.

We test these predictions in a preregistered experiment with 560 US consumers. We contribute to research on consumer responses to AI (Puntoni, Reczek, Giesler, & Botti, 2021) by showing that AI disclosure is not uniformly costly, and to research on authenticity (Morhart, Malär, Guèvremont, Girardin, & Grohmann, 2015; Newman & Dhar, 2014) by identifying AI authorship as a new threat to it.

Conceptual Background

Consumer Responses to AI

A large literature documents that consumers often prefer human to algorithmic judgment. Dietvorst et al. (2015) showed that people abandon algorithms after seeing them make mistakes, even when the algorithm outperforms humans. Longoni et al. (2019) found that consumers resist medical AI because they believe it cannot account for their unique characteristics. Preferences for human labor are especially strong in symbolic consumption contexts (Granulo, Fuchs, & Puntoni, 2021), and products labelled as handmade are seen as containing "love" (Fuchs, Schreier, & van Osselaer, 2015). Recent work shows that AI-generated marketing content is perceived as less authentic than human-written content (Bergstrøm & Lind, 2023).

Hedonic Versus Utilitarian Products

Hedonic products are consumed primarily for affective and sensory experience, utilitarian products primarily for their function (Dhar & Wertenbroch, 2000). Because hedonic consumption is about experience and meaning, we expect consumers to care more about the human origin of the product and its story. A description written by AI breaks the link between the product and a human who cared about it. For utilitarian products, the description is mainly a carrier of information, and its source should matter less.

H1: An AI-disclosure label lowers purchase intention.

H2: The negative effect of an AI-disclosure label on purchase intention is stronger for hedonic than for utilitarian products.

The Role of Authenticity

Authenticity, the perception that something is genuine and true to its origin, is a central driver of brand and product evaluations (Morhart et al., 2015). Authenticity can transfer from the original source of production (Newman & Dhar, 2014). We argue that AI authorship undermines this transfer for hedonic products: an AI-written description signals that no one in particular stands behind the product.

H3: Perceived authenticity mediates the effect of the AI-disclosure label on purchase intention for hedonic products.

Study 1

Method

Participants. We recruited 560 US participants on Prolific (Peer, Brandimarte, Samat, & Acquisti, 2017) in exchange for a small payment. Following our preregistration (AsPredicted #118804), we excluded participants who failed the attention check (n = 33; Oppenheimer, Meyvis, & Davidenko, 2009). We also excluded participants who completed the survey in under two minutes (n = 45) and participants who reported a willingness to pay above $100 (n = 6). The final sample consisted of 476 participants (M_{age} = 37.42, SD = 11.48; 51.7% women). [OLA: should we say more about Prolific here? Reviewers at EMAC asked.]

Design and procedure. The study used a 2 (label: AI-written vs. no label) × 2 (product: scented candle vs. phone charger) between-subjects design. Participants were randomly assigned to one of the four conditions and read a short product description on a mock online store page. In the AI-label conditions, a small badge below the description read "This description was written by AI." The descriptions were identical across label conditions and matched in length across products.

After reading the description, participants indicated whether the description had been written by AI (yes, by AI / no / not sure; manipulation check). They then reported their purchase intention, perceived authenticity, and willingness to pay, completed an attention check, and provided demographics.

Measures. Purchase intention was measured with three items (e.g., "How likely would you be to buy this product?"; 1 = very unlikely, 7 = very likely; α = .91). Perceived authenticity was measured with four items adapted from Morhart et al. (2015) (e.g., "This product feels genuine"; 1 = strongly disagree, 7 = strongly agree; α = .56). Willingness to pay was an open-ended question ("What is the most you would pay for this product, in US dollars?").

Results

Manipulation check. More participants in the AI-label conditions than in the no-label conditions reported that the description had been written by AI (77.7% vs. 22.6%; $\chi^2(1) = 144.19$, $p < .001$). The manipulation was successful.

Purchase intention. Table 2 shows the means by condition. As predicted, the AI label lowered purchase intention for the scented candle (M = 4.16 vs. 4.50; $t(225) = -2.37$, $p = .018$, $d = -0.32$). For the phone charger, the difference was not significant (M = 4.79 vs. 4.94; $t(247) = -1.01$, $p = .314$, $d = -0.13$). Hence, product type moderated the effect of the AI label: the label hurt the hedonic product but not the utilitarian one, supporting H2. Across products, the label lowered purchase intention, in line with H1.

Mediation. To test H3, we ran a mediation analysis for the candle using PROCESS Model 4 (Hayes, 2018) with 5,000 bootstrap samples. The AI label lowered perceived authenticity (a = −0.49), and authenticity increased purchase intention (b = 0.43). The indirect effect was significant (indirect effect = −0.21, 95% CI [−0.34, -0.10]), while the direct effect was not ($c' =$

−0.13). Perceived authenticity thus explains why AI disclosure lowers purchase intention for hedonic products, supporting H3.[1]

General Discussion

An AI-disclosure label lowered purchase intention for a hedonic product but not for a utilitarian one, and perceived authenticity explained the effect. These findings speak to a timely managerial question, and they suggest that the costs of transparency about AI are not uniform.

Theoretical contributions. First, we extend research on algorithm aversion (Castelo et al., 2019; Dietvorst et al., 2015) from algorithmic advice to algorithmic marketing content. Second, we identify AI authorship as a threat to perceived authenticity, adding to work on the sources of authenticity (Newman & Dhar, 2014). Third, we show that the hedonic–utilitarian distinction (Dhar & Wertenbroch, 2000) predicts when AI disclosure is costly.

Managerial implications. Sellers of hedonic products should be cautious about AI-written descriptions, or should accompany the disclosure with cues of human involvement ("written with AI, checked by our founder"). Sellers of utilitarian products can likely use AI copywriting with little risk, even when they disclose it.

Limitations and future research. Our study used one hedonic and one utilitarian product and a US online sample. Future research should examine other product categories, real purchase behaviour, and whether cues of human involvement can offset the effect. [KN: Study 2 (field data from the Etsy-like platform) goes here once Ola has the data.]

References

Bergstrøm, H., & Lind, M. (2023). Machine-made meaning: Authenticity inferences from AI-generated marketing content. *Journal of Consumer Psychology, 33*(4), 512–529. https://doi.org/10.1002/jcpy.1399

Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research, 56*(5), 809–825. https://doi.org/10.1177/0022243719851788

Dhar, R., & Wertenbroch, K. (2000). Consumer choice between hedonic and utilitarian goods. *Journal of Marketing Research, 37*(1), 60–71. https://doi.org/10.1509/jmkr.37.1.60.18718

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144*(1), 114–126. https://doi.org/10.1037/xge0000033

Fuchs, C., Schreier, M., & van Osselaer, S. M. J. (2015). The handmade effect: What's love got to do with it? *Journal of Marketing, 79*(2), 98–110. https://doi.org/10.1509/jm.14.0018

Granulo, A., Fuchs, C., & Puntoni, S. (2021). Preference for human (vs. robotic) labor is stronger in symbolic consumption contexts. *Journal of Consumer Psychology, 31*(1), 72–80. https://doi.org/10.1002/jcpy.1181

Hayes, A. F. (2018). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (2nd ed.). Guilford Press.

[1] Willingness to pay showed a directionally consistent pattern for the candle (M = $14.50 vs. $15.60), but the difference was not significant ($t(221) = -1.13$, $p = .259$). Because the open-ended willingness-to-pay measure is noisy, we focus on purchase intention.

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes, 151*, 90–103. https://doi.org/10.1016/j.obhdp.2018.12.005

Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research, 46*(4), 629–650. https://doi.org/10.1093/jcr/ucz013

Morhart, F., Malär, L., Guèvremont, A., Girardin, F., & Grohmann, B. (2015). Brand authenticity: An integrative framework and measurement scale. *Journal of Consumer Psychology, 25*(2), 200–218. https://doi.org/10.1016/j.jcps.2014.11.006

Newman, G. E., & Dhar, R. (2014). Authenticity is contagious: Brand essence and the original source of production. *Journal of Marketing Research, 51*(3), 371–386. https://doi.org/10.1509/jmr.11.0022

Nordmann, K., & Nordmann, O. (2024). Who wrote this? Source cues and quality inferences in product descriptions. In *Proceedings of the European Marketing Academy Conference (EMAC)*.

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology, 45*(4), 867–872. https://doi.org/10.1016/j.jesp.2009.03.009

Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology, 70*, 153–163. https://doi.org/10.1016/j.jesp.2017.01.006

Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing, 85*(1), 131–151. https://doi.org/10.1177/0022242920953847

Tables

Table 1. Sample.

	N	Age, M (SD)	Women (%)
Final sample	476	37.42 (11.48)	51.7

Table 2. Means (standard deviations) by condition.

Product	Label	N	Purchase intention	Perceived authenticity
Scented candle	AI-written	106	4.16 (1.09)	3.80 (0.80)
Scented candle	No label	121	4.50 (1.07)	4.29 (0.86)
Phone charger	AI-written	127	4.79 (1.16)	4.26 (0.80)
Phone charger	No label	122	4.94 (1.14)	4.38 (0.77)