

When Gender Meets Super-Expertise: Perceptual Bias and Behavioral Effects in AI–Human Hybrid Recommendations

MARCO FRANCESCO MAZZÙ¹, LORENZO RICCIARDI¹, MATTEO DE ANGELIS¹

Abstract. Consumers increasingly rely on algorithmic and human recommendations to navigate complex purchase decisions, yet research has predominantly examined either human or AI recommendation sources in isolation. This study focuses on AI–human hybrid recommendation systems, in which algorithmic outputs are combined with human expertise. Specifically, it examines whether recommender gender and expertise level jointly shape consumer perceptions and behavioral responses. Using a 2×2 between-subjects experiment ($N = 169$), participants evaluated an AI–human hybrid recommendation for a virtual reality headset. Gender was treated as a social cue and expertise as a boundary condition. Results show a differentiated pattern across expertise conditions. When the recommender is described as an expert, perceived transparency does not differ significantly between male and female recommenders. When the recommender is described as a super-expert, the female recommender is perceived as more transparent than the male recommender. In addition, perceived transparency was negatively associated with purchase intention, indicating that greater informational openness may activate caution rather than commitment in high-technology experience contexts. The paper provides preliminary evidence relevant to the debate on AI–human hybrid systems, suggesting that gender and expertise may function as joint determinants of perceived transparency, and offering initial insights into the conditions under which social identity cues shape consumer responses within algorithmic recommendation environments.

Keywords: Product recommendations; Artificial Intelligence; AI–human hybrid systems; Expertise signaling; Gender bias; Experience goods; Behavioral decision-making.

Framing of the research.

AI-assisted recommendations in experience-based, high-technology contexts. AI-assisted recommendation systems are increasingly designed as hybrid configurations, in which algorithmic outputs are combined with human expertise to guide consumer decisions (Puntoni et al., 2021; Mazzù et al., 2025). Examples of these configurations include e-commerce platforms that combine algorithmic product recommendations with advice from certified specialists, or retail interfaces that integrate AI-generated suggestions with human advisors presented as industry experts. This design logic rests on the assumption that human cues — such as credentialed expertise and recognizable identity — can enhance the legitimacy of algorithmic advice and reduce the uncertainty consumers face when evaluating complex or unfamiliar products. Yet the conditions under which human source characteristics shape how hybrid recommendations are perceived and acted upon remain poorly understood.

Existing research has primarily compared human and AI recommendation sources, while largely overlooking how identity cues embedded in the human component of hybrid systems influence consumer evaluations. Do social identity cues such as gender influence how transparent an AI–human recommendation system is perceived to be? Does this effect depend on the level of expertise attributed to the human recommender? And does any perceptual influence translate into purchase intention? These questions are especially relevant for experience-based, high-technology products, where objective pre-purchase evaluation is limited and consumers must rely heavily on external signals to guide economically consequential decisions.

In hybrid environments, the distinction between AI-generated and human-provided advice is increasingly blurred. Algorithmic systems supply analytical depth and scale; human experts provide contextual framing, legitimacy, and socio-emotional resonance (Puntoni et al., 2021; De Cicco et al., 2025). Prior work has established that recommendation source type shapes consumer reactions (Mazzù et al., 2024), but substantially less attention has been devoted to the identity characteristics of the human component embedded within hybrid systems. In high-technology experience contexts, where product quality cannot be fully assessed before purchase (Nelson, 1970) and financial stakes are salient, the characteristics of the human source — including their perceived expertise and their gender — may function as diagnostic cues that shape how the entire recommendation system is evaluated. Specifically, such cues may influence perceived transparency (Mazzù et al., 2023), defined here as perceived informational openness — the extent to which the recommendation is seen as providing clear, accessible, and candid product information — rather than algorithmic transparency concerning how the AI system operates.

This paper examines precisely this question. Focusing on gender as a social cue and expertise level as a boundary condition, we test whether their interaction shapes perceived transparency of an AI-human hybrid recommendation and whether perceived transparency, in turn, predicts purchase intention. The study extends prior work on source effects in recommendation systems (Baccelloni et al., 2025a) by moving beyond binary human-vs-AI comparisons to examine how the identity characteristics of the human component moderate perceptual evaluations within a constant AI-supported framework. By doing so, it addresses a gap in research that has predominantly focused on attitudinal outcomes while

¹ LUISS Business School, Via Nomentana 216, 00162 Roma, LUISS Guido Carli University, Department of Business and Management, Viale Romania 32, 00197 Roma, Italia

leaving the process mechanisms and behavioral consequences of human source cues in hybrid AI environments underexplored.

Gender and perceptual legitimacy in AI-assisted high-technology contexts. *Gender is among the most pervasive social categorization cues in interpersonal and mediated communication (Eagly & Karau, 2002). In source evaluation contexts, gender operates not as an explicit evaluative criterion but as an implicit heuristic that activates gender-role expectations and stereotypic associations. In domains stereotypically linked to technical authority — such as engineering, computing, and high-technology consumption — these associations often favor male sources. When a recommendation is presented by a male versus female expert, consumers may draw on gendered competence schemas to interpret the information (Kim et al., 2024; Calluso et al., 2025).*

Research on gender and expertise evaluations consistently shows that the influence of gender depends on contextual moderators, including task type and perceived congruence between the domain and gender-role expectations (Calluso et al., 2025; Eagly & Karau, 2002). In technology-related contexts, male sources are generally perceived as more congruent with the domain, which may lead female sources — particularly those positioned at high levels of authority — to face evaluative penalties linked to expectancy violation (Rudman & Glick, 2001). Importantly, this effect is most likely to emerge when the recommender's identity is made salient — as occurs when expertise is strongly signaled. At the same time, expectancy violations may also produce the opposite effect. When individuals occupy positions that contradict stereotypical expectations, their behavior and communication style may be evaluated more carefully and perceived as more deliberate and informative. In such cases, a female recommender presented as a super-expert may be perceived as providing more transparent and detailed information precisely because her authority signals exceptional competence within a domain where expectations are typically male-oriented. In conditions of lower expertise, when identity cues are less diagnostic, gender effects on perceptual evaluations may be suppressed or absent.

While prior work has examined gender effects on trust and credibility perceptions in advisory contexts, less attention has been devoted to how gender shapes perceived transparency — an evaluative dimension particularly relevant to AI-assisted settings, where consumers assess the extent to which the recommendation provides clear guidance. Perceived transparency in hybrid recommendation systems is not simply a property of the algorithmic component; it is also colored by the identity of the human source through whom the recommendation is communicated. If a female super-expert is perceived as occupying a role at odds with domain expectations, the informational openness of the entire recommendation may be interpreted more carefully, even when the content of the advice is identical.

A further consideration concerns the link between perceived transparency and purchase behavior. In high-technology, experience-based contexts, greater informational openness does not necessarily translate into greater commitment. On the contrary, detailed disclosure of product complexity and limitations may increase consumers' awareness of risk and uncertainty, leading to more cautious allocation decisions (Schilke & Reimann, 2025). This “transparency caution” mechanism provides an important theoretical complement to the perceptual effects of source cues: even if gender and expertise shape how transparent the recommendation is perceived to be, the downstream behavioral consequence of transparency may reflect risk activation rather than trust formation.

Expertise as a legitimacy cue and boundary condition under uncertainty. *A foundational insight from source credibility research is that, when consumers cannot directly evaluate the quality of information, they rely on source characteristics as proxies for reliability and relevance (Hovland et al., 1953; Chaiken, 1980; Pornpitakpan, 2004). Expertise — the degree of domain-relevant knowledge and experience attributed to a source — is among the most consistently studied of these cues. In complex, technically sophisticated product categories, high expertise is generally associated with greater persuasiveness and stronger recommendation uptake (Donovan et al., 2024). Expertise signals not only technical competence but also normative authority, indicating that the recommender is institutionally qualified to guide decisions under uncertainty (Wiener & Mowen, 1986; Pornpitakpan, 2004).*

In AI-assisted recommendation environments, the authority function of expertise becomes particularly salient. When consumers encounter an AI-human hybrid system, the presence of a credentialed human expert may serve as a signal of informational legitimacy — not merely of the human's personal knowledge, but of the entire recommendation configuration. This is consistent with dual-process accounts, which suggest that source cues become especially influential under conditions of uncertainty or when cognitive elaboration is constrained, leading consumers to rely on heuristic processing rather than systematic evaluation (Petty & Cacioppo, 1986). Empirical research in AI-assisted retail contexts extends this principle, showing that expertise-related cues modulate how recommendation systems are evaluated as a whole, with higher expertise signals increasing the perceived authority of the overall human-algorithm configuration (Chinchanchokchai et al., 2021). For high-technology experience products, where most consumers lack the domain expertise to independently assess recommendation quality, the perceived authority of the human source may thus shape how the system as a whole is evaluated.

Importantly, elevated expertise may also amplify the diagnostic value of other source characteristics. When a recommender is perceived as a routine expert, consumers may focus primarily on the informational content of the recommendation. When the recommender is positioned as an outstanding authority — for instance, through extensive experience, certifications, and institutional recognition — the identity of the source itself becomes more salient. Under these conditions, expertise functions not only as a direct credibility cue but as a boundary condition that modulates how additional characteristics, including gender, are interpreted and weighed. This amplification logic suggests that the effect

of gender on perceptual evaluations may emerge specifically in high-expertise conditions, where the recommender's identity is most foregrounded.

This perspective treats expertise level as a structural feature of the recommendation environment rather than a simple additive predictor. Its relevance lies not in whether more expertise leads to better outcomes, but in whether it changes the interpretive context within which other cues — such as gender — operate. We accordingly treat expertise as a boundary condition in the analytical model.

Purpose of the paper. This paper examines how human source cues embedded in AI-assisted recommendation systems shape consumer perceptions and purchase intention. Specifically, it investigates (i) whether recommender gender and expertise jointly influence perceived transparency in an AI-human hybrid recommendation, and (ii) whether perceived transparency, in turn, predicts purchase intention in a high-technology experience context. Gender is treated as a social cue, expertise as a boundary condition, perceived transparency as the focal perceptual outcome, and purchase intention as the downstream behavioral response. Rather than assuming a fully supported mediation chain, the study assesses whether social identity cues operate primarily at the perceptual stage and whether their effects extend downstream to behavioral outcomes. Based on the foregoing theoretical arguments, the following hypotheses are advanced:

H1. The effect of recommender gender on perceived transparency of an AI-human hybrid recommendation is moderated by expertise level, such that gender differences in transparency perceptions are more pronounced when the recommender is described as a super-expert than when described as an expert.

H2. Perceived transparency of an AI-human hybrid recommendation negatively predicts purchase intention, such that higher levels of perceived transparency are associated with lower purchase intention toward the recommended product.

RQ (Exploratory). As an exploratory question, this study examines whether the indirect effect of recommender gender on purchase intention through perceived transparency differs as a function of expertise level (i.e., whether the conditional indirect effect is consistent with a moderated mediation pattern). Given the partial and marginal nature of the first-stage interaction (H1), no directional hypothesis is advanced regarding the full indirect path.

Methodology. The study employed a 2 (human recommender gender: Female = 0 vs. Male = 1) $\times$ 2 (expertise level: expert = 0 vs. super-expert = 1) between-subjects experimental design within an AI-assisted recommendation context. Participants were randomly assigned to one of four conditions and exposed to an online shopping scenario involving a virtual reality (VR) headset, selected as a high-technology experience product characterized by high ex ante uncertainty. Across conditions, recommendations were consistently presented as AI-assisted, with an identifiable human expert integrated into the system; only the recommender's gender and level of expertise were manipulated. Expertise was manipulated by describing the recommender either as an industry expert or as a highly credentialed authority with over 30 years of experience, while all other elements of the stimulus (product information, layout, and wording) were held constant. As intended, the expertise manipulation was successful: participants in the super-expert condition perceived the recommender as significantly more expert than those in the expert condition, confirming the effectiveness of the manipulation (see Appendix).

Descriptive statistics indicated that perceived transparency averaged 4.34 ($SD = 1.36$), and purchase intention equaled 2.01 ($SD = .46$). Recommender gender was approximately balanced across conditions ($M = 0.51$, $SD = 0.50$), as was expertise level ($M = 0.51$, $SD = 0.50$). Bivariate correlations revealed a negative association between perceived transparency and purchase intention ($r = -.378$). Gender and expertise level showed negligible correlations with purchase intention ($r = .153$ and $r = -.001$, respectively) and with each other ($r = .065$).

After exposure to the recommendation scenario, participants completed a behavioral allocation task in which they were asked to distribute a hypothetical budget of \$499 — matching the price of the VR headset — across three mutually exclusive options: purchase the recommended product (fully following the recommendation), seek additional information (partial deferral), or retain the budget (full rejection of the recommendation). This three-option forced allocation design, scored on a 1–3 scale, was chosen as the primary dependent variable because it captures financial commitment as a consequential and mutually exclusive decision rather than as a self-reported attitude. Unlike intention scales, the budget allocation task imposes a concrete trade-off and thereby introduces behavioral realism into an experimental setting (Chandon et al., 2005). The task is particularly appropriate for experience products involving financial risk, where the willingness to allocate real resources constitutes a meaningful proxy for downstream purchase behavior. Response distributions indicated that 10.1% of participants chose to purchase (option 1), 78.7% opted to seek additional information (option 2), and 11.2% retained the budget (option 3), confirming that the distribution across response categories, while concentrated on option 2, reflects a behaviorally coherent pattern in a high-uncertainty, high-involvement product context. While the compressed 1–3 range may reduce statistical sensitivity relative to continuous scales — a limitation acknowledged in the Research Limitations section — the forced-choice structure is theoretically justified by the discrete and mutually exclusive nature of the financial commitment it represents. Although the budget was hypothetical, prior research has shown that incentive-compatible framing of trade-off tasks elicits behaviorally consistent responses in experimental consumer research. Participants also completed self-reported measures of intention to follow the recommendation; results on this variable have been reported elsewhere (Mazzù et al., 2025) and serve here as a reference for situating the behavioral findings.

Additional perceptual measures assessed perceived transparency, credibility, warmth, competence, and gender fit using validated multi-item 7-point Likert-type scales adapted from prior research. Perceived transparency (5 items; e.g.,

“I believe the recommendation source will provide me with transparent product information”) was operationalized as perceived informational openness — the extent to which the recommendation source was seen as providing clear, accessible, and in-depth product information — rather than procedural or algorithmic transparency per se, following Mazzù et al. (2023). Credibility was measured with 10 items assessing both trustworthiness (e.g., dependable, honest, reliable) and expertise (e.g., expert, knowledgeable, qualified) dimensions, adapted from Hovland et al. (1953) and Pornpitakpan (2004). Warmth and competence were each assessed with 3-item scales following established social perception research. Gender fit was assessed through two complementary 6-item scales capturing both general domain-gender congruence beliefs and the perceived suitability of the specific recommender for the product category. All scales were drawn from validated instruments. These measures were included to support the interpretation of behavioral outcomes; however, in line with the conditional process framework, perceived transparency was specified as the focal mediator in the analytical model.

Data were collected via Prolific from a European sample. Of 224 initial respondents, 169 participants (aged 18–65) passed attention checks and were retained for analysis, yielding approximately 42 participants per cell. Participants were randomly assigned to one of the four experimental conditions. A post-hoc power analysis using G*Power (Faul et al., 2007) for a multiple regression model with three predictors and a medium effect size ($f^2 = .15$, $\alpha = .05$) indicated that a sample of 169 provides statistical power of approximately .90. For the moderated mediation pathway, in which effect sizes tend to be smaller, the present sample provides sufficient power for an initial test of the proposed framework, though future studies should employ larger samples to detect more distal conditional indirect effects with greater precision.

Effects of the experimental manipulations were examined using Hayes’ (2022) PROCESS macro for SPSS (Model 7), testing a first-stage moderated mediation model in which recommender gender predicted purchase intention indirectly through perceived transparency, conditional on expertise level. Interaction terms were computed directly from these coded variables. Indirect effects were estimated using 5,000 bootstrapped samples with 95% confidence intervals. All analyses were conducted using SPSS.

Results.


Interaction effects on perceived transparency. The regression model predicting perceived transparency explained a modest portion of variance ($R^2 = .026$, $F(3,165) = 1.48$, $p = .221$). Neither the main effect of recommender gender ($b = -0.19$, $p = .372$) nor the main effect of expertise ($b = 0.12$, $p = .571$) was statistically significant. However, the interaction between gender and expertise was marginally significant ($b = -0.77$, $SE = .42$, $p = .067$), indicating that the effect of gender on perceived transparency depends on the level of expertise attributed to the recommender.

Conditional effects analysis showed that when the recommender was portrayed as a super-expert, gender negatively predicted perceived transparency ($b = -0.56$, $p = .056$), whereas no significant gender effect emerged in the expert condition ($b = 0.21$, $p = .486$). This pattern indicates that gender differences in perceived transparency arise only when expertise is framed as particularly high. Specifically, in the expert condition male recommenders were perceived as slightly more transparent than female recommenders, whereas in the super-expert condition the opposite pattern emerged. Figure 1 illustrates this interaction.

The indirect effect of gender on purchase intention through perceived transparency was not statistically significant at either level of expertise. In the expert condition, the indirect effect was small and non-significant, and in the super-expert condition it was larger in magnitude but also non-significant (both 95% CIs included zero). Consistent with these results, the index of moderated mediation was not statistically significant (index = .097, 95% CI $[-.008, .212]$).

Overall, the results indicate that gender and expertise jointly influence early perceptual evaluations of transparency, but these perceptual differences do not translate into statistically reliable conditional indirect effects on purchase intention. The findings therefore provide partial support for the proposed interaction hypothesis, suggesting that expertise acts as a boundary condition shaping how gender cues are interpreted at the perceptual stage.

Figure 1. Moderated mediation model – estimated coefficients



Summary of results. Overall, the results provide partial support for the proposed model. The interaction between gender and expertise on perceived transparency was marginally significant, with gender differences emerging only in the super-expert condition. Perceived transparency negatively predicted purchase intention, suggesting that greater informational openness may activate caution rather than commitment in a high-uncertainty technology context. The conditional indirect effects of gender on purchase intention through transparency were not statistically significant at either level of expertise, and the index of moderated mediation was likewise non-significant. Taken together, these findings suggest that human source cues in AI-assisted recommendation environments operate primarily at the early perceptual stage, with limited downstream effects on behavioral outcomes.

Research limitations. This study has several limitations that suggest directions for future research. First, the analysis focuses on a single high-technology, experience-based product category. Although a virtual reality headset is well suited to examining decision-making under uncertainty, future research should assess whether similar dynamics emerge across other experience goods and in mixed search–experience contexts. Second, the sample size is sufficient for detecting the focal direct effects but remains relatively limited for moderated mediation analysis, where interaction-based indirect effects tend to be small in magnitude and require larger samples to achieve adequate statistical power. The marginal significance of the first-stage interaction ($p = .067$) and the non-significant indirect effects may in part reflect this power constraint rather than a true absence of the hypothesized processes. Future research should therefore employ larger samples specifically designed to provide adequate power for detecting conditional indirect effects of the expected magnitude. Third, the study concentrates on gender and expertise as the focal human cues embedded within AI-assisted recommendations. Other characteristics—such as age, communication style, perceived similarity, or visual presentation—may shape consumer responses in different ways. Future studies should investigate how broader combinations of human attributes influence perceptual mechanisms in hybrid AI–human systems. Fourth, the dependent variable — a three-option forced allocation task scored on a 1–3 ordinal scale — while theoretically grounded in behavioral realism, imposes constraints on statistical sensitivity. The compressed response range may reduce variance and limit power to detect small conditional effects, particularly in the moderated mediation pathway. Future research might explore continuous allocation formats (e.g., percentage-based budget distribution from 0 to 100%) or multi-item behavioral indices to improve measurement resolution. Fifth, although the budget allocation task required participants to make a concrete and mutually exclusive monetary trade-off, it remained hypothetical. Research on hypothetical bias suggests that stated willingness to pay or act may diverge from actual behavior, particularly when financial stakes are salient (Chandon et al., 2005). Future research might employ incentive-compatible designs — such as bonus-conditional choice tasks on crowdsourcing platforms — or field experiments in partnership with retail operators. Finally, further theoretical work is needed to clarify the role of perceived transparency in high-uncertainty contexts. The present findings suggest that informational openness may not always facilitate commitment and, under some conditions, may instead activate caution. Future studies should examine more closely when transparency enhances trust and when it heightens perceived complexity or risk, thereby refining its role in AI-assisted decision environments.

Managerial implications. For managers, the findings suggest that evaluating AI-assisted recommendation systems solely through attitudinal metrics may provide an incomplete picture of system performance. Although the direct behavioral effects of recommender gender were limited, the results indicate that human cues can shape perceptual interpretations of transparency under specific expertise conditions. This highlights the importance of complementing attitudinal indicators with behavioral metrics when assessing the effectiveness of hybrid AI–human recommendation systems. The findings also indicate that integrating highly credentialed human experts into AI-supported interfaces is not a neutral design choice. Elevated expertise increases the salience of the recommender’s identity, which may activate gender-related interpretive biases in technology-related contexts. Managers should therefore consider not only the credentials they communicate, but also how the identity of the human recommender interacts with expertise framing in shaping users’ early evaluations of the system.

In addition, the negative association between perceived transparency and purchase intention suggests that informational openness does not always function as a straightforward conversion mechanism. In high-technology, experience-based contexts, detailed disclosure may increase awareness of product complexity and uncertainty. Managers should therefore carefully consider how and when transparency-related information is presented along the customer journey.

For platform designers, the way human experts are visually and descriptively integrated into AI-supported interfaces may influence how the recommendation system as a whole is interpreted by users. Interface elements such as displayed credentials, biographies, or authority signals can increase the salience of identity cues and shape how recommendations are evaluated.

Finally, the results highlight the importance of the moment of financial commitment within AI-assisted customer journeys. Interface design efforts aimed at supporting clarity, reducing cognitive overload, and managing perceived uncertainty at the point of action may be more consequential than further optimizing upstream trust perceptions alone.

Originality of the paper. This study contributes to research on AI-human hybrid recommendation systems by examining whether social identity cues embedded in the human component of such systems influence early-stage perceptual evaluations and whether these evaluations relate to purchase intention. The results provide partial evidence that gender

differences in perceived transparency emerge when the human recommender holds a very high level of signaled expertise — a finding consistent with the view that elevated authority amplifies the salience of identity-related cues in uncertain, high-technology evaluation contexts. Theoretically, the study contributes at the intersection of expectancy violation theory (Rudman & Glick, 2001) and source credibility research (Hovland et al., 1953; Pornpitakpan, 2004) by showing that the evaluative consequences of role-incongruent identity cues are not uniform but are contingent on the intensity of authority signaling. In particular, the finding that a female super-expert — rather than a standard expert — elicits higher perceived transparency aligns with expectancy violation accounts: when a communicator occupies a role that defies stereotypical expectations and does so at an especially high level of authority, the resulting incongruence may focus evaluative attention on the informational content rather than activating backlash. This mechanism extends prior work by specifying expertise level as the structural condition under which gender-based expectancy violations shift from producing evaluative penalties to producing perceptual amplification. Furthermore, by connecting this perceptual amplification to a “transparency caution” effect on purchase behavior — consistent with Schilke and Reimann’s (2025) transparency dilemma framework — the paper integrates source credibility, social identity, and behavioral decision-making literatures within a unified process model of hybrid AI–human recommendation influence.

Perceived transparency negatively predicted purchase intention, suggesting that informational openness may activate risk awareness rather than facilitate commitment in technology-uncertainty contexts — a finding that qualifies commonly held assumptions about transparency as a uniformly trust-enhancing mechanism. The conditional indirect effects linking gender to purchase intention through transparency were not statistically significant, which the study interprets as evidence of attenuation along the causal chain rather than as disconfirmation of the perceptual-stage effects.

More broadly, the study contributes to a process-oriented understanding of AI-human hybrid influence in digital retail by clarifying the scope and limits of human source cues within recommendation environments. The results indicate that social identity signals may shape perceptual stage evaluations even within algorithmic systems, but that their downstream behavioral consequences are more limited and context-dependent than prior attitudinal research might suggest. This boundary-condition perspective offers a grounded basis for further research on when and how hybrid AI-human designs translate source characteristics into consumer behavior.

In sum, the paper offers a calibrated account of hybrid recommendation influence: human source characteristics can shape how algorithmic advice is perceptually interpreted, but this influence may attenuate before reaching economically consequential outcomes. Transparently reporting where the evidence is strong and where it remains partial represents, in itself, a contribution to the growing body of research seeking to understand the conditions under which hybrid AI-human systems exert behavioral effects on consumers.

References

- ADOMAVICIUS, G., TUZHILIN, A. (2005), “Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions”, *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, n. 6, pp. 734-749.
- BACCELLONI, A., DE ANGELIS, M., RICOTTA, F., MAZZÙ, M.F. (2025), “Too Narrow to Help? Unveiling How Recommendation Agents’ Specialization Impacts User Choices”, *Journal of Interactive Marketing*.
- BACCELLONI, A., MAZZÙ, M.F., RICOTTA, F., MATTIACCI, A. (2025), “Uncovering the role of weak ties in implicit networks of influence: a network analysis on recommendation algorithms’ neighborhood”, *European Journal of Marketing*, vol. 59, n. 12, pp. 2725-2762.
- CALLUSO, C., et al. (2025), “Gender, leadership, and cultural context”, *Proceedings of the WOA 2025 Conference*.
- CHAIKEN, S. (1980), “Heuristic versus systematic information processing and the use of source versus message cues in persuasion”, *Journal of Personality and Social Psychology*, vol. 39, n. 5, pp. 752-766.
- DE CICCO, R., FRANCONI, B., DINI, M., CURINA, I., FILIERI, R., CIOPPI, M. (2025), “Generative AI and Human Advice: User Perceptions, Intentions, and Behavior Across Pre-visit and On-site Travel Stages”, *Journal of Travel Research*.
- CHINCHANACHOKCHAI, S., DUFF, B.R.L., MANCHANDA, P. (2021), “A tale of two recommender systems: The moderating role of consumer expertise on artificial intelligence based product recommendations”, *Journal of Retailing and Consumer Services*, vol. 61, 102528.
- DIETVORST, B.J., SIMMONS, J.P., MASSEY C. (2015), “Algorithm aversion: People erroneously avoid algorithms after seeing them err”, *Journal of Experimental Psychology: General*, vol. 144, n. 1, pp. 114-126.
- DONOVAN, R.J., JALLEH, G., DRANE, C. (2024), “Impact of source credibility on behavioural responses to a mental health promotion social marketing campaign”, *Journal of Social Marketing*, vol. 14, n. 2, pp. 250-263.
- HOVLAND, C.I., JANIS, I.L., KELLEY, H.H. (1953), *Communication and persuasion*, Yale University Press, New Haven.
- KARMARKAR, U.R., TORMALA, Z.L. (2010), “Believe me, I have no idea what I’m talking about: The effects of source certainty on consumer involvement and persuasion”, *Journal of Consumer Research*, vol. 36, n. 6, pp. 1033-1049.
- KIM, J.H., SONG, H. (2020), “The influence of perceived credibility on purchase intention via competence and authenticity”, *International Journal of Hospitality Management*, vol. 90.
- KIM, S., OH, P., LEE, J. (2024), “Algorithmic gender bias: investigating perceptions of discrimination in automated decision-making”, *Behaviour & Information Technology*, vol. 43, n. 16, pp. 4208-4221.
- LONGONI, C., BONEZZI, A., MOREWEDGE, C.K. (2019), “Resistance to medical artificial intelligence”, *Journal of Consumer Research*, vol. 46, n. 4, pp. 629-650.
- MAZZÙ, M.F., POZHARLIEV, R., ANDRIA, A., & BACCELLONI, A. (2023). Overcoming the blockchain technology credibility gap. *Psychology & Marketing*, 40(9), 1791-1807.
- MAZZÙ, M.F., DE ANGELIS, M., ANDRIA, A., BACCELLONI, A. (2024), “Humans or AI? How the source of recommendations

- influences consumer choices for different product types”, *California Management Review*.
- MAZZÙ, M.F., ANDRIA, A., BACCELLONI, A., et al. (2025), “The impact of humans vs. AI recommendation on consumer reactions to products exposure”, *Electronic Commerce Research*.
- NELSON, P. (1970), “Information and consumer behavior”, *Journal of Political Economy*, vol. 78, n. 2, pp. 311-329.
- ÖZBÖLÜK, T., AKDOĞAN, K. (2022), “The role of online source credibility and influencer identification on consumers’ purchase decisions”, *International Journal of Internet Marketing and Advertising*, vol. 16, n. 1, pp. 1-18.
- PETTY, R.E., CACIOPPO, J.T. (1986), *Communication and persuasion: Central and peripheral routes to attitude change*, Springer-Verlag, New York.
- PORNPITAKPAN, C. (2004), “The persuasiveness of source credibility: A critical review of five decades' evidence”, *Journal of Applied Social Psychology*, vol. 34, n. 2, pp. 243–281.
- PUNTONI, S., RECZEK, R.W., GIESLER, M., BOTTI, S. (2021), “Consumers and artificial intelligence: An experiential perspective”, *Journal of Marketing*, vol. 85, n. 1, pp. 131-151.
- SCHILKE, O., REIMANN, M. (2025), “The transparency dilemma: How AI disclosure erodes trust”, *Organizational Behavior and Human Decision Processes*, vol. 188.
- WIENER, J.L., MOWEN, J.C. (1986), “Source Credibility: On The Independent Effects Of Trust And Expertise”, *ACR North American Advances*.
- CHANDON, P., MORWITZ, V.G., REINARTZ, W.J. (2005), “Do Intentions Really Predict Behavior? Self-Generated Validity Effects in Survey Research”, *Journal of Marketing*, vol. 69, n. 2, pp. 1–14.
- EAGLY, A.H., KARAU, S.J. (2002), “Role Congruity Theory of Prejudice Toward Female Leaders”, *Psychological Review*, vol. 109, n. 3, pp. 573–598.
- FAUL, F., ERDFELDER, E., LANG, A.G., BUCHNER, A. (2007), “G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences”, *Behavior Research Methods*, vol. 39, n. 2, pp. 175–191.
- RUDMAN, L.A., GLICK, P. (2001), “Prescriptive Gender Stereotypes and Backlash Toward Agentic Women”, *Journal of Social Issues*, vol. 57, n. 4, pp. 743–762.

Appendix

Figure A1. AI-assisted recommendation with male expert




Figure A2. AI-assisted recommendation with female expert




Figure A3. AI-assisted recommendation with male super-expert

Home

Products

About us

FAQ

Login



VIRTUAL REALITY HEADSET

Recommended based on cutting-edge AI application combined with insights from an industry expert with more than 30 years of experience in the field



Year: 2022

Color: Black

Technical Specifications:

- Cutting-edge technology: yes
- Processing power: high level
- Screen: 5.7-inch OLED screen
- Immersive experience: High
- Image: Images at 120 frames per second
- Latency: Minimal
- Cooperative gaming capability: Yes
- Audio: 3D audio
- High-performance multimedia engine
- Design: Refurbished, compact and cylindrical

...

Figure A4. AI-assisted recommendation with female super-expert

Home

Products

About us

FAQ

Login



VIRTUAL REALITY HEADSET

Recommended based on cutting-edge AI application combined with insights from an industry female expert with more than 30 years of experience in the field



Year: 2022

Color: Black

Technical Specifications:

- Cutting-edge technology: yes
- Processing power: high level
- Screen: 5.7-inch OLED screen
- Immersive experience: High
- Image: Images at 120 frames per second
- Latency: Minimal
- Cooperative gaming capability: Yes
- Audio: 3D audio
- High-performance multimedia engine
- Design: Refurbished, compact and cylindrical

...