



Article

Strategic Disclosure of AI Curation: A Boundary Condition on Algorithm Aversion in Hedonic E-Commerce

Tiannv Ma ^{1,†}, Yuqi Du ^{2,†}, Yong Wang ^{3,†} and Liying Zhou ^{4,*}

¹ School of Management, Changchun University, Changchun 130022, China; matn@ccu.edu.cn

² International Chinese College, Panyapiwat Institute of Management, Nonthaburi 11120, Thailand; dyq@mail.gufe.edu.cn

³ College of Business and Public Management, West Chester University, West Chester, PA 19383, USA; ywang2@wcupa.edu

⁴ Guian New Area Institute, Guizhou University of Finance and Economics, Guiyang 550025, China

* Correspondence: zhouly@mail.gufe.edu.cn

† These authors contributed equally to this work and share first authorship.

Abstract

Algorithm-aversion research predicts that consumers prefer human curators to algorithmic ones in subjective decision domains, including taste-based hedonic recommendation. Drawing on the algorithmic-symbiosis paradigm and on assortment-perception theory, this paper identifies a boundary condition on that prediction: disclosing that recommendations are AI-curated rather than human-curated lifts purchase intention in hedonic e-commerce but not in utilitarian e-commerce. The mechanism is a search-side option-breadth inference—the consumer’s attribution about the size of the option pool the curator considered upstream—which is diagnostic in preference-formative consumption categories where consumers build, rather than match, a preference. Three online experiments deployed through a Chinese consumer panel test the framework. Study 1 ($N = 228$) finds the predicted Disclosure $\times$ Product-type interaction ($\eta_p^2 = 0.025$) with the AI-versus-human lift confined to the hedonic cell ($d = 0.65$). Study 2 ($N = 257$) isolates the option-breadth pathway against trust and competence as competing mediators. Study 3 ($N = 519$) extends the design to a second hedonic category, decomposes option breadth into search-side and display-side subdimensions through an eight-item bi-factor scale, tests mentalizing alongside option breadth as a competing mediator, and brings the moderated-mediation test by consumer AI familiarity to conventional statistical power (index of moderated mediation = +0.065, 95% CI [+0.014, +0.118]). Implications for interactive-marketing practice and algorithmic-disclosure regulation are discussed.

Keywords: algorithmic symbiosis; AI disclosure; algorithm aversion; perceived option breadth; hedonic e-commerce; recommendation agents; interactive marketing



Academic Editor: Albert Y.S. Lam

Received: 28 June 2026

Accepted: 16 July 2026

Published: 18 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

E-commerce platforms now openly badge their recommendations as the work of machines. Streaming services advertise “AI-DJ” playlists and “AI-Picks” carousels, content marketplaces watermark feeds as “For You” engines, and Chinese hedonic platforms such as NetEase Cloud Music surface algorithmically chosen tracks under prominent neural-network logos [1,2]. The visibility of these disclosures is striking because the dominant theoretical prediction in the literature would counsel against them. A decade of algorithm-aversion research has documented that decision-makers resist algorithms in subjective

or taste-based domains, which are the very domains in which hedonic e-commerce operates [3–6]. The empirical practice of platforms therefore sits in tension with the empirical practice of academia: brands advertise AI loudly where the literature predicts they should hide it. The unresolved question is whether platforms are reading their consumers more accurately than the literature has, and, if so, why. The stakes of resolving the tension have been raised further by an unfolding mandatory-disclosure regulatory regime—China’s Cyberspace Administration of China [7,8] Algorithm Recommendation Provisions and Interim Measures for Generative AI Services, and the European Union [9] Article 50 transparency obligation—which gives the academic question direct policy traction.

This paper identifies a boundary condition that qualifies the convergent algorithm-aversion prediction in hedonic e-commerce. The argument is that the prediction “subjective domains beget aversion” depends on what inference the consumer draws about the algorithmic source. When consumers infer that an AI has computationally traversed a wider option pool than any human curator could feasibly traverse, the disclosure stops functioning as a signal of impersonal coldness and starts functioning as a signal of curatorial breadth. The shift is consequential because perceived option breadth is itself a well-established positive driver of consumer satisfaction, choice confidence, and willingness to engage [10,11]. The deeper reason the inference becomes diagnostic in hedonic categories is that hedonic consumption is *preference-formative*: the consumer arrives at the decision without a fully specified ideal point and uses the choice process to construct the preference [12–15]. Under preference-formative conditions, larger considered pools deliver discovery value rather than diagnostic noise. In utilitarian categories, where consumers arrive with a specification (a price ceiling, a feature threshold), the same inferred breadth is rationally indifferent once specification is met.

Three online experiments test the framework. Study 1, a two-by-two between-subjects design ($N = 228$), crosses curator disclosure (AI vs. human) with product type (hedonic vs. utilitarian) and tests the predicted interaction on purchase intention. Study 2 ($N = 257$) holds the hedonic context constant, formally tests the option-breadth mediation pathway against competing trust and competence pathways, and examines whether consumer AI familiarity moderates the indirect effect. Study 3 ($N = 519$) extends the design to a second hedonic category for stimulus generalisability, decomposes option breadth into search-side and display-side subdimensions through an eight-item bi-factor scale, tests mentalizing alongside option breadth as a competing mediator, and brings the moderated-mediation test by consumer AI familiarity to conventional statistical power.

The paper develops three contributions. The first contribution is a falsifiable boundary condition on the algorithm-aversion-in-subjective-domains pattern that has anchored a decade of consumer–AI research [3,6,16]. The qualification is conditional rather than universal: the hedonic-cell AI-versus-human lift survives across two product categories and two studies (S1 + S3), while the utilitarian-cell null also reproduces (S1). The second contribution is a mediator (search-side perceived option breadth) that is conceptually and empirically distinct from competence, warmth, generalized trust, anthropomorphism, and mentalizing—the five mediators that have so far carried the AI-disclosure effect in published work. The construct draws on the assortment-perception tradition [11] and operates through the preference-formative nature of hedonic consumption rather than through attributions about the curator as a social agent. The third contribution is a regulatory and managerial framing: mandatory algorithm-disclosure regulation, designed under a friction-creates-attenuation assumption, may interact heterogeneously with category structure. In hedonic categories, the same disclosure cue regulators reason will dampen persuasion can amplify it through the search-side breadth pathway. This paper closes with

a concrete decision rule for platforms operating hedonic surfaces and a set of bounded implications for the algorithmic-symbiosis agenda in interactive marketing.

The remainder of this paper develops the theoretical argument and the hypotheses (Section 2), reports Study 1 (Section 3) and Study 2 (Section 4), provides a cross-study bridge (Section 5), reports Study 3 (Section 6), and discusses theoretical, regulatory, and managerial implications (Section 7) before concluding (Section 8).

2. Theoretical Background and Hypotheses

2.1. Algorithm Aversion and Appreciation: A Convergence and Its Discontents

The dominant theoretical claim about consumer–algorithm interaction has, for nearly a decade, been algorithm aversion. Dietvorst, Simmons, and Massey [4] showed that decision-makers abandon objectively superior algorithms after observing a single forecasting error, which is an asymmetric reaction not produced by equivalent errors from human forecasters. The pattern generalized quickly across managerial estimation [17], medical recommendations [5], and dietary advice [18]. Two mechanisms recur in the explanations offered: a perceived inability of algorithms to accommodate particularised cases, which Longoni, Bonezzi, and Morewedge (2019) [5] called uniqueness neglect, and a felt loss of consumer agency when human judgement is overruled by code [19]. A meta-analytic review of twenty-nine studies codified both the pattern and its mechanisms [20]. The most influential boundary condition on the aversion contest is Castelo, Bos, and Lehmann's (2019) [3] task-subjectivity moderator: people delegate to algorithms when tasks feel objective and resist when tasks feel subjective. Wien and Peluso [6] extended the moderator directly to retail recommendation, finding that human recommenders outperform AI recommenders on hedonic products because human recommenders elicit stronger mentalizing responses that support consumer self-referencing.

A counter-current runs in parallel. Logg, Minson, and Moore [21] surfaced an algorithm-appreciation effect in which participants gave more weight to algorithmic advice than to human advice for forecast-class tasks particularly when their own expertise was low. Bogert, Schechter, and Watson [22] extended the result to harder tasks, in which difficulty itself triggered increased reliance on algorithmic input. A recent meta-analysis of more than four hundred effect sizes [16] reconciles the two streams within a capability–personalization framework: AI appreciation occurs when the AI is perceived as more capable than humans *and* personalization is not required, while AI aversion intensifies whenever personalization is perceived as necessary. The framework provides a high-altitude organizing scheme in which both empirical patterns find their place.

The convergent prediction from the aversion lineage for hedonic e-commerce is unambiguous: an AI-disclosure cue should suppress consumer responses relative to a human-disclosure cue. We argue that this prediction holds only under one specific framing of the AI source—namely, as an impersonal, socially distant judge whose evaluations stand in opposition to human warmth or human mentalizing. Once the AI source is framed as a *curator* in a category where the consumer's preference is being formed rather than confirmed, a competing inference becomes diagnostic: an algorithm has searched through a larger option pool than any human could feasibly traverse, and that reach matters in a category where consumers expect to discover variety. We develop this preference-formation argument in Section 2.2 and connect it to the search-side option-breadth mediator in Section 2.3 before stating the hypotheses in Sections 2.4–2.6. Table 1 summarizes the position of this paper within the prior literature.

Table 1. Prior studies of AI-disclosed vs. human-disclosed recommendation effects: a positioning map.

Study	Context	Disclosure Manipulation	Direction (AI vs. Human)	Mediator	How Our Paper Differs
Dietvorst, Simmons, and Massey [4]	Forecasting	Algorithm vs. human	Aversion after error	Error asymmetry	Recommendation rather than forecasting; appreciation in preference-formative categories
Castelo, Bos, and Lehmann [3]	Various; experimental	Algorithm vs. human	Aversion for subjective tasks	Subjectivity perception	Boundary within subjective tasks (hedonic e-commerce); option-breadth mediator
Logg, Minson, and Moore [21]	Forecast tasks	Algorithm vs. crowd advice	Appreciation	Capability inference	Appreciation extended to a domain where aversion has been predicted
Longoni, Bonezzi, and Morewedge [5]	Medical AI	Algorithm vs. human physician	Aversion	Uniqueness neglect	Consumer e-commerce; predicts appreciation in hedonic
Yeomans, Shah, Mullainathan, and Kleinberg [18]	Recommendation reasoning	Algorithm vs. human recommender	Mixed; sense-making route	Inferred reasoning	Inferred upstream pool rather than reasoning route
Bogert, Schecter, and Watson [22]	Difficulty estimation	Algorithm vs. human advisor	Appreciation, esp. under difficulty	Self-uncertainty	Preference uncertainty in hedonic categories
Wien and Peluso [6]	Travel-product recommendation	Algorithm vs. human travel agent	Aversion in hedonic	Mentalizing	Measures mentalizing alongside option breadth in Study 3; both pathways present, option-breadth dominant in net effect
Longoni and Cian [23]	Hedonic vs. utilitarian goods	Disclosure as AI-generated	Hedonic–utilitarian asymmetry	Affect/“word-of-machine”	Distinct mechanism (search-side option breadth) for the asymmetry
Cian, Longoni, and Krishna [24]	Domain-unfamiliar AI use	Disclosure as AI-generated	Lift under unfamiliarity	Domain knowledge	Within-domain inference about curator reach
Qin et al. [16]	Meta-analytic synthesis	Various	Capability × personalization reconciliation	—	Instantiates the framework in retail recommendation with a new mediator
<i>Present paper</i>	Hedonic vs. utilitarian e-commerce	AI-curated vs. human-curated badge	Hedonic-cell lift; utilitarian-cell null	Search-side perceived option breadth	Three experiments; bi-factor mediator decomposition; competing-mediator test against mentalizing; AI-familiarity moderation at conventional power

2.2. Preference-Formative vs. Preference-Confirming Consumption

The convergent aversion prediction for hedonic e-commerce treats hedonic and utilitarian consumption as if they differed only in the affective valence of the product. We argue that the more consequential distinction lies in the cognitive structure of the consumer’s decision. Hedonic consumption is *preference-formative*: the consumer arrives at the decision without a fully specified ideal point and uses the choice process itself to construct the preference. Utilitarian consumption is *preference-confirming*: the consumer arrives with a specification (a price ceiling, a feature threshold, a category specification) and uses the choice process to match. The distinction has a downstream consequence for what consumers infer when a curator pre-filters the choice set.

The cognitive-science backdrop for the formative claim is preference construction. Slovic [15] argued, in a synthesis that has since shaped a generation of behavioral-decision research, that preferences are constructed rather than retrieved; the decision-maker is not a passive utility maximizer but an active builder of evaluative criteria during the choice process. Bettman, Luce, and Payne [25] brought the constructive-choice frame-

work into consumer research, demonstrating that the choice process itself shapes the preference being expressed. Within hedonic consumption specifically, the constructive pattern is sharper. Holbrook and Hirschman [26] introduced the experiential view of consumption—fantasy, feeling, fun—in which the experience *is* the product and the consumer’s evaluation criteria emerge from interaction with the option set. Hoch [13] develops the seductive-product mechanism by which consumers learn what they want only after exposure; Botti and Iyengar [12] document the choice-difficulty asymmetry between hedonic preference uncertainty and utilitarian specification matching; Khan, Dhar, and Wertenbroch [14] provide a behavioral-decision-theory synthesis of the formative–confirming distinction at the choice-process level.

The asymmetry has a clean implication for assortment perception. Under preference-formative conditions, larger considered pools deliver *discovery* value: the consumer learns more about her own preference by being shown more variety, particularly variety she would not have generated through search alone. Chernev [27] develops the ideal-point-availability mechanism by which assortments yield value differently depending on whether the consumer has a pre-formed ideal point; pre-formative choosers benefit from larger assortments, post-formative choosers do not. Chernev, Böckenholt, and Goodman’s (2015) [28] meta-analysis of choice overload reinforces the point: the demotivation effect of large assortments documented by Iyengar and Lepper [29] attenuates to non-significance when a curator pre-filters the set, which is an outcome that is consistent with the curator absorbing the discovery burden that the preference-formative consumer would otherwise carry alone.

For the disclosure question, the formative–confirming distinction shifts the consumer’s inference about the curator’s upstream input set. An AI-disclosed curator activates a schema in which the curator has computationally searched a substantially larger option pool than a human curator could feasibly traverse—a schema consistent with the contemporary technology backdrop in which algorithmic systems are widely understood as capable of processing very large datasets. The schema-driven inference is not a technical claim about actual algorithmic reach; it is the folk-theoretic content the disclosure cue activates in the consumer’s mind. In preference-formative hedonic categories, that activation is converted into a discovery payoff and ultimately into a higher purchase intention. In preference-confirming utilitarian categories, the same activation occurs at the inference level but is rationally indifferent at the purchase-intention level because the consumer’s evaluative criterion is specification matching rather than discovery. The disclosure cue therefore produces an interaction signature on purchase intention—positive in hedonic categories, null in utilitarian ones—that the next three subsections formalize as the focal mediation pathway (H1–H2), its job-design parallels and competitor mechanisms, and the consumer-level moderator (H3) that scales the indirect effect.

2.3. Search-Side Perceived Option Breadth: Definition and Discriminant Validity

We define **search-side perceived option breadth** as the consumer’s inference about the size of the option pool the curator considered, attributed to the curator’s search effort, distinct from the displayed variety on the recommendation surface itself. The construct is therefore an attribution about an unobserved upstream search rather than a count of observed downstream items. Two implications follow. The construct can vary while the visible recommendation set is held constant, which is precisely the design feature exploited in our stimuli, where the four cells share a pixel-identical visible set and differ only in a single curator-disclosure badge. And the construct attaches inferentially to the curator, so curator-identity cues can shift it without changing the assortment displayed.

The construct must be distinguished from a small constellation of neighboring concepts that the assortment-perception, choice-uncertainty, and human–AI interaction that have

been separately developed in the literature. *Perceived variety* [11] refers to the within-set heterogeneity of a displayed assortment; *objective assortment size* [27] refers to the count of options in that displayed set. Neither construct addresses the upstream pool the curator drew from. *Recommendation quality* [30] is a fit judgement between the recommendation and the consumer's own preference—a judgement that, by the argument in Section 2.2, the preference-formative hedonic consumer cannot reliably render because her preference is still being constructed through the choice process. *Platform capability* or *AI competence* [31] is a trait-level inference about the agent's general technical capacity, which is distinct from the behavior-level inference about the agent's behavior in the present instance; the near-zero zero-order correlation between option breadth and AI competence observed in our Study 2 ($r = +0.04$, $p = 0.50$, $N = 257$) is consistent with this conceptual separation. *Mentalizing* [6,32] is an inference about whether the agent has a mind, which is rooted in the social-cognition tradition rather than the assortment-perception tradition that motivates option breadth.

Within our own framing, we further decompose option breadth into a *search-side* subdimension (inference about the upstream considered pool) and a *display-side* subdimension (perception of the visible recommendation surface). The decomposition is theoretically motivated by the Section 2.2 argument: the search-side inference is what carries the discovery payoff in preference-formative hedonic categories, and the display-side perception is a related but secondary input. The two subdimensions are operationalized through an eight-item bi-factor scale developed for Study 3 with confirmatory factor analysis used to test whether the bifurcation is supported by the data. The scale used in Studies 1 and 2 mixed search-side and display-side language at the item level; we acknowledge this as a methodological limitation of those earlier studies and use the bi-factor scale to refine the construct measurement in this paper's main empirical contribution.

The distinction between search-side option breadth and mentalizing deserves explicit empirical treatment because mentalizing is the closest competing mediator in the AI-disclosure literature. Three lines of argument anchor the discriminant claim. First, the two constructs have different referents: mentalizing is about whether the agent has a mind, search-side option breadth is about how large the agent's input pool was. A consumer can hold a high search-side option-breadth inference without attributing any mind to the AI (for example, "the algorithm scraped half a million products" is a search-side inference without mind attribution). Second, the two constructs have different correlational profiles: if option breadth were a mentalizing variant, it would correlate non-trivially with trust or competence (both of which load on mind perception); the near-zero correlations observed in Study 2 are empirical support for distinctness. Third, the two constructs have different theoretical antecedents: mentalizing derives from social-cognition theory [32]; search-side option breadth derives from assortment-perception theory [10,11]. The two types of literature have grown independently, and the constructs were conceived without reference to each other. Study 3 provides the head-to-head competing-mediator test.

2.4. H1: AI-Disclosure $\times$ Product-Type Interaction on Purchase Intention

The framework developed in Section 2.2 implies a specific empirical signature: an interaction between curator-disclosure and product type on consumer purchase intention. The argument runs in six linked steps. First, hedonic consumption is preference-formative—the consumer arrives at the decision without a fully specified ideal point, and the choice process itself constructs the preference [12,13,15,25]. Second, under preference-formative conditions, the consumer's payoff from a curated set scales with the inferred breadth of the considered pool—not with the displayed assortment alone—because larger considered pools deliver discovery, and discovery is the experiential payoff hedonic consumers

seek [27,28]. Third, utilitarian consumption is preference-confirming: the consumer arrives with a specification, and once the displayed set satisfies the specification, the inferred breadth of the upstream considered pool is rationally indifferent to the consumer's purchase decision [14]. Fourth, an AI-disclosure cue activates a schema in which the curator is taken to have considered a substantially larger option pool than a human curator would. Fifth, combining the second and fourth steps, in hedonic categories, an AI-disclosure cue produces a larger search-side option-breadth inference, which is converted into a larger discovery payoff, which is in turn expressed as a higher purchase intention. Sixth, combining the third and fourth steps, in utilitarian categories, the same disclosure cue still shifts the search-side inference, but purchase intention is invariant to that inference once specification is met.

H1. *Curator-disclosure (AI vs. human) interacts with product type (hedonic vs. utilitarian) in determining consumer purchase intention.*

H1a. *In hedonic product categories, AI-disclosed recommendations elicit higher purchase intention than human-disclosed recommendations.*

H1b. *In utilitarian product categories, AI-disclosed and human-disclosed recommendations elicit purchase intention indistinguishable in mean.*

The hypothesis is framed as a *boundary condition* on, rather than a reversal of, the algorithm-aversion convergence. The qualification is conditional: the hedonic-cell positive contrast and the utilitarian-cell null contrast are both required; the empirical signature is the interaction term rather than either cell-contrast alone.

2.5. H2: Search-Side Perceived Option Breadth Mediates the Hedonic-Cell Effect

If H1 holds, the hedonic-cell positive contrast requires a mechanism. The Section 2.2 chain identifies search-side perceived option breadth as the proximal mediator: the AI-disclosure cue activates the larger-search-pool inference, the inference is converted into a discovery payoff in the preference-formative hedonic context, and the discovery payoff drives purchase intention. The mediation is distinct from the four mediators that have so far carried the AI-disclosure effect in published work—competence and warmth [31,33], generalized trust in the recommendation [34], perceived anthropomorphism [24,35], and mentalizing [6]—because the four prior mediators all draw on attributions about the curator as a social actor, while search-side option breadth draws on an attribution about the curator's computational reach.

H2. *In hedonic categories, the effect of AI-disclosure on purchase intention is mediated by search-side perceived option breadth.*

The bi-factor decomposition of perceived option breadth into search-side and display-side subdimensions, operationalized in Study 3, generates an empirical refinement of H2. The argument in Section 2.2 predicts that the search-side subdimension does most of the explanatory work in hedonic categories, because it is the upstream input pool—not the visible recommendation surface—that delivers the discovery payoff under preference formation. The display-side subdimension is expected to contribute a smaller, positive pathway through the more familiar within-set variety mechanism [11].

H2a. *When perceived option breadth is decomposed into search-side and display-side subdimensions, the mediation effect is carried primarily by the search-side subdimension; the display-side subdimension carries a smaller, positive indirect effect.*

2.6. H3: Consumer AI Familiarity Moderates the Mediated Effect

A natural question is whether the H2 mediation is uniform across consumers or whether it scales with consumer characteristics. We argue that consumer AI familiarity, defined as the consumer's prior exposure to and self-rated knowledge of AI-enabled products, moderates the search-side mediation through a four-step chain anchored in the capability–personalization framework and persuasion-knowledge theory.

First, the broader algorithm-aversion and appreciation literature exhibits an unresolved tension with both directions supported by substantial empirical bases [3,21,22]. The Qin et al. [16] meta-analysis reconciles the tension within a capability–personalization framework: consumers prefer algorithmic advice when the algorithm is perceived as capable and when personalization is not required. AI familiarity is the consumer-level antecedent of the capability dimension of this framework: familiar consumers carry a richer schema of what AI can do, which calibrates upward the capability perception triggered by an AI-disclosure cue [36,37].

Second, AI familiarity is also the consumer-level antecedent of a reduced persuasion-knowledge defensiveness response. The persuasion-knowledge model [38] predicts that consumers who recognize a persuasion attempt activate defensive coping that dampens the persuasive effect. Consumers who interact regularly with AI-driven recommendations in everyday digital life—streaming services, shopping platforms, content feeds—do not interpret an AI-disclosure cue as a fresh persuasion attempt to defend against; their default schema is informational rather than adversarial. Recent meta-analytic evidence places consumer experience at the center of the AI-acceptance picture [39], which is consistent with this mechanism.

Third, the two consumer-level shifts compound. Higher capability perception of the AI curator, combined with lower persuasion-knowledge defensiveness, produces a larger inference about how much the AI curator considered—a larger search-side option-breadth inference. Familiar consumers extend more schema credit to the AI-disclosure cue.

Fourth, since H2 establishes search-side option breadth as the proximal mediator of the disclosure effect on purchase intention, the consumer-level amplification of the search-side inference at higher AI familiarity translates directly into a stronger conditional indirect effect.

H3. *Consumer AI familiarity moderates the search-side perceived option breadth mediation such that the conditional indirect effect of AI-disclosure on purchase intention via search-side perceived option breadth is stronger at higher levels of consumer AI familiarity.*

3. Study 1

3.1. Method

Adult online shoppers were recruited through Credamo, which is a leading Chinese consumer panel widely used in marketing research. The advertised sample size was 260, which was calibrated against an a priori power analysis (G*Power 3.1; two-by-two interaction, $\alpha = 0.05$, power = 0.80, anticipated $f = 0.18$) that returned a required N of approximately 240. After applying a priori exclusion rules (failed attention checks, mismatched manipulation-check responses, or completion times outside the 60 to 1500 s window), the analytic sample comprised 228 respondents (51 percent female; $M_{\text{age}} = 33.0$ years). All inferential tests are two-tailed at $\alpha = 0.05$. We treat the H1 hedonic-

cell AI-versus-human contrast and the H2 search-side option-breadth mediation as focal tests specified in advance; secondary contrasts are reported descriptively without family-wise correction. The study was approved by the Ethics Committee of Guian New Area Institute, Guizhou University of Finance and Economics (Approval Code GNARIRB0013, 15 October 2025).

The design crossed curator disclosure (AI vs. human) with product type (hedonic vs. utilitarian) in a fully between-subjects allocation. Stimuli were custom-built mock-up product pages: a music-subscription playlist landing page for the hedonic category and a wireless-mouse listing for the utilitarian category. The two cells of each product paired pixel-identical screenshots that differed only in a single 14-character curation-label badge. The full visual specifications of all four cells appear in Appendix B.

Participants gave informed consent, completed brief demographics, and were randomly assigned to one of the four stimulus screenshots. A 30-s arithmetic distractor separated the stimulus from the dependent measures to clear short-term memory effects. Respondents then completed a forced-choice curator-disclosure manipulation check, the Voss, Spangenberg, and Grohmann [40] hedonic/utilitarian manipulation check, the four-item perceived-option-breadth scale developed for this paper (single-factor form), the four-item Komiak and Benbasat [34] trust scale (item 4 was reworded from “competent” to “credible” to avoid lexical overlap with the AI-competence construct), the three-item Dodds, Monroe, and Grewal [41] purchase-intention scale, the three-item MacKenzie and Lutz [42] attitude-toward-recommendation semantic differential, the three-item Fiske, Cuddy, and Glick [31] AI-competence scale, the four-item Tian, Bearden, and Hunter [43] need-for-uniqueness short form, and two attention checks. Scale blocks were presented in a fixed order across respondents rather than randomized; we acknowledge this as a methodological limitation in Section 7 and note that the Study 3 design corrects it through block-order randomization. The median completion time was 6.2 min. All multi-item Likert scales used a seven-point format; composite scores are arithmetic means of items. The full-item wording is reproduced verbatim in Appendix A. The novel perceived-option-breadth scale was pretested in a separate 30-respondent pilot to validate item clarity through cognitive interviews [44] and then submitted to a three-expert content-validity panel that confirmed a scale-level CVI of 0.92.

A confirmatory factor analysis of the four perceived-option-breadth items at Study 1 ($N = 228$) returned acceptable fit, $\chi^2/df = 2.34$, CFI = 0.98, RMSEA = 0.07, with standardized loadings ranging from 0.82 to 0.91, AVE = 0.73, and composite reliability CR = 0.92. Heterotrait–monotrait ratios were computed from a three-construct CFA model containing perceived option breadth, trust in recommendation, and perceived AI competence: HTMT against trust = 0.71 and against perceived AI competence = 0.68, which were both below the conventional 0.85 cutoff [45], supporting discriminant validity. The same set of analyses at Study 2 ($N = 257$) returned $\chi^2/df = 2.18$, CFI = 0.97, RMSEA = 0.07, loadings 0.80 to 0.89, AVE = 0.70, CR = 0.90, and HTMT values of 0.73 (trust) and 0.69 (competence) from the corresponding three-construct CFA. We note that the single-factor scale used in Studies 1 and 2 mixes search-side and display-side language at the item level, which is a measurement issue that was addressed in the refined bi-factor scale developed for Study 3.

3.2. Results

All retained respondents correctly identified the assigned curator source on the forced-choice manipulation check, which is a mechanical consequence of the exclusion criterion that removed disclosure-MC failures. The hedonic stimulus scored higher on the hedonic semantic-differential than on the utilitarian semantic-differential, and the utilitarian stimu-

lus showed the opposite pattern; both manipulation effects were significant at conventional levels, confirming the product-type manipulation. All composite measures cleared the conventional $\alpha \geq 0.80$ reliability threshold (Table 2).

Table 2. Reliability and descriptive statistics, Study 1 ($N = 228$).

Construct	<i>k</i> Items	<i>M</i>	<i>SD</i>	α
Purchase intention	3	5.085	1.019	0.885
Perceived option breadth	4	5.154	1.091	0.925
Trust in recommendation	4	4.995	1.025	0.905
AI competence	3	5.078	1.019	0.876
Consumer AI familiarity	4	4.737	1.372	0.950
Hedonic [40]	5	4.512	1.412	0.963
Utilitarian [40]	5	5.026	1.215	0.957
Need for uniqueness	4	4.182	0.684	0.852

Note: Means, standard deviations, and Cronbach's α are reported to three decimal places to disambiguate constructs whose rounded two-decimal summaries (purchase intention and AI competence) happen to coincide.

The four cell means for purchase intention are consistent with the predicted interaction. AI-disclosed recommendations elicited higher purchase intention than human-disclosed recommendations in the hedonic cell, while the two utilitarian cells were indistinguishable (Table 3; Figure 1).

Table 3. Purchase intention cell means, Study 1.

Product Type	Disclosure	<i>M</i>	<i>SD</i>	<i>n</i>
Hedonic	AI-curated	5.47	1.06	55
Hedonic	Human-curated	4.83	0.88	59
Utilitarian	AI-curated	5.03	1.05	54
Utilitarian	Human-curated	5.03	1.00	60




Figure 1. Disclosure $\times$ Product type interaction on purchase intention (Study 1).

A two-by-two analysis of variance on purchase intention returned a significant Disclosure $\times$ Product type interaction, $F(1,224) = 5.81$, $p = 0.017$, partial $\eta^2 = 0.025$, 90% CI [0.003, 0.063] [46], with non-significant main effects for both disclosure ($F < 1$) and product type ($F = 1.23$, $p = 0.27$). Decomposing the interaction, AI-disclosure produced significantly higher purchase intention than human-disclosure within the hedonic stimulus ($M_{AI} = 5.47$, $SD = 1.06$; $M_{Human} = 4.83$, $SD = 0.88$), $t(112) = 3.46$, $p < 0.001$ (two-tailed), Cohen's $d = 0.65$. The same contrast within the utilitarian stimulus returned a null effect ($M_{AI} = 5.03$; $M_{Human} = 5.03$), $t(112) = -0.01$, $p = 0.99$. H1a is supported in the predicted direction; H1b is supported in the null sense predicted. Adding need for uniqueness as a covariate left the interaction term essentially unchanged, $F(1,223) = 5.78$, $p = 0.017$. Adding perceived AI competence as a covariate attenuated but did not eliminate the interaction, $F(1,223) = 4.39$, $p = 0.037$, suggesting that the lift is not absorbed by competence inferences. A post hoc power calculation against the observed interaction effect ($f = 0.16$, computed from $\eta_p^2 = 0.025$) at $\alpha = 0.05$ with $N = 228$ returns observed power ≈ 0.67 , which we report transparently as below the 0.80 conventional target; the higher-powered confirmation lies in Study 3. We re-estimated the focal interaction on the full pre-exclusion sample ($N = 260$, no attention-check or manipulation-check exclusions) and on an attention-pass-only intermediate sample ($N = 246$). The interaction survives in both: $F(1,256) = 4.09$, $p = 0.044$, partial $\eta^2 = 0.016$ in the full sample, and $F(1,242) = 5.02$, $p = 0.026$, partial $\eta^2 = 0.020$ in the attention-pass-only sample. The H1 pattern is therefore not an artefact of the manipulation-check-driven exclusion rule.

3.3. Discussion

Study 1 produces the predicted Disclosure $\times$ Product-type interaction on purchase intention. The hedonic-cell AI-versus-human contrast is positive and significant; the utilitarian-cell contrast is null. The finding sits opposite the prediction issued by the Castelo–Wien–Peluso lineage within the hedonic cell while reproducing the algorithm-appreciation null in the utilitarian cell. The mechanism question—why does the AI-disclosure cue lift purchase intention specifically in the hedonic cell?—is taken up in Study 2 through a formal mediation test, and it is deepened in Study 3 through a bi-factor decomposition of the mediator and a head-to-head competition with mentalizing as an alternative pathway.

4. Study 2

4.1. Method

A separate sample of 300 adult online shoppers was recruited through Credamo with disclosure cells balanced one-to-one-to-one across an AI condition, a human condition, and a no-label baseline. The a priori exclusion rules used in Study 1 retained 257 respondents (53 percent female; $M_{age} = 32.6$ years). Power for the mediation test was calibrated using Fritz and MacKinnon's (2007) [47] tables for medium-sized paths; we acknowledge in advance that the moderated-mediation test (H3) was underpowered at this sample size with conventional power recovered in Study 3. As in Study 1, all inferential tests are two-tailed at $\alpha = 0.05$, and the H2 indirect effect is treated as the focal test. The study was approved under the same IRB protocol as Study 1.

The hedonic stimulus from Study 1 (music-subscription playlist landing page) was held constant. The AI and human cells used the same screenshots as Study 1, and the no-label baseline removed the curator badge entirely while leaving every other visual element pixel-identical. The no-label cell serves as a descriptive baseline against which to calibrate the AI and human cells; it is not part of the focal H2 mediation, which is theoretically defined as the comparison between AI-attributed and human-attributed curation. Consumer

AI familiarity was measured before, rather than after, stimulus exposure to avoid post-treatment contamination. All other measures, exclusion rules, and procedural details mirrored Study 1 with the limitation that mentalizing was not fielded in the Study 2 instrument. This omission, which we address in the competing-mediator test in Study 3, is acknowledged transparently in Section 4.3.

4.2. Results

After exclusions, all retained participants correctly identified their assigned disclosure cell. Internal consistencies again cleared the conventional threshold (Table 4).

Table 4. Reliability and descriptive statistics, Study 2 ($N = 257$).

Construct	<i>k</i> Items	<i>M</i>	<i>SD</i>	α
Purchase intention	3	5.05	0.94	0.85
Perceived option breadth	4	4.92	0.97	0.90
Trust in recommendation	4	4.93	1.03	0.91
AI competence	3	5.06	0.95	0.86
Consumer AI familiarity	4	4.80	1.27	0.95

Note: Zero-order correlations among the three candidate mediators in Study 2 are: perceived option breadth $\times$ trust $r = -0.04$, $p = 0.54$; perceived option breadth $\times$ AI competence $r = 0.04$, $p = 0.50$; trust $\times$ AI competence $r = 0.08$, $p = 0.19$ (all $N = 257$).

The three-cell pattern broadly tracks the Study 1 finding: AI-disclosed recommendations produce the highest mean purchase intention, human-disclosed recommendations sit slightly below, and the no-label baseline sits lowest (Table 5). The differential is particularly pronounced for perceived option breadth, where the AI cell sits well above both alternatives, which is consistent with the breadth-signal account that motivates H2.

Table 5. Mean responses by disclosure cell, Study 2.

Disclosure	<i>M</i> (PI)	<i>SD</i>	<i>M</i> (OB)	<i>SD</i>	<i>M</i> (TR)	<i>SD</i>	<i>n</i>
AI-disclosed	5.22	0.82	5.44	0.86	5.04	0.90	86
Human-disclosed	5.03	0.92	4.72	0.95	4.89	0.98	84
No label	4.93	1.04	4.66	0.97	4.78	1.15	87

A one-way ANOVA on purchase intention across the three cells was marginal at the omnibus level, $F(2, 254) = 2.47$, $p = 0.087$ (two-tailed). The omnibus F is not the focal test; the focal test is the AI-versus-human pairwise contrast that carries the H2 mediation. Using PROCESS Model 4 [48] with X coded as AI-versus-human, M as perceived option breadth, and Y as purchase intention, the a -path was robustly positive ($a = +0.61$, $p < 0.001$), the b -path positive and significant ($b = +0.52$, $p < 0.001$), the total effect $c = +0.20$, and the direct effect c' dropped to -0.12 after controlling for option breadth (Table 6). The bootstrapped indirect effect ab was $+0.32$ with a 95-percentile percentile bootstrap confidence interval of $[+0.16, +0.49]$ (5000 resamples), excluding zero. The opposite-sign direct effect ($c' = -0.12$) relative to the total effect ($c = +0.20$) is a classic case of inconsistent mediation [49], indicating that the perceived-option-breadth pathway more than absorbs the disclosure cue's net effect on purchase intention. H2 is supported. Re-estimating the mediation with trust in recommendation as the mediator yielded a non-significant indirect effect ($ab = +0.08$, 95% percentile bootstrap CI $[-0.04, +0.21]$); a parallel re-estimation with perceived AI competence as the mediator also returned a non-significant indirect effect ($ab = +0.09$, 95% percentile bootstrap CI $[-0.03, +0.22]$). Perceived option breadth thus outperforms both a generalized trust-based account and a competence-based account, reinforcing the claim that the mediator is mechanism-distinct from the social-actor pathways

that dominate the existing AI-disclosure literature. The Study 2 mediator race did not include a mentalizing measure of the kind Wien and Peluso [6] identify as the mechanism in their hedonic-recommendation studies; the mentalizing scale was not fielded in the Study 2 instrument, and Study 3 closes this gap with a head-to-head test.

Table 6. Mediation results, Study 2 (X = AI vs. Human, M = perceived option breadth, Y = purchase intention).

Path	Estimate	95% Percentile Bootstrap CI
<i>a</i> (Disclosure → Option breadth)	+0.61	—
<i>b</i> (Option breadth → Purchase intention)	+0.52	—
<i>c'</i> (Direct effect Disclosure → PI)	−0.12	—
<i>c</i> (Total effect Disclosure → PI)	+0.20	—
<i>ab</i> (Indirect effect)	+0.32	[+0.16, +0.49]
Index of Moderated Mediation (<i>W</i> = consumer AI familiarity)	+0.045	[−0.07, +0.16]

To probe H3, the mediation was re-estimated with PROCESS Model 7 allowing consumer AI familiarity (mean-centered) to moderate the *a*-path. The Index of Moderated Mediation was +0.045 with a 95-percent percentile bootstrap confidence interval of [−0.07, +0.16]. The point estimate is in the predicted direction, but its confidence interval contains zero, so at this sample size, H3 is directionally supported but does not clear the conventional significance threshold. Conditional indirect effects rose monotonically across the consumer-AI-familiarity distribution: +0.27, 95% percentile bootstrap CI [+0.06, +0.51] at −1 SD; +0.33, 95% CI [+0.17, +0.50] at the mean; and +0.39, 95% CI [+0.18, +0.62] at +1 SD. Each conditional indirect effect excludes zero on its own CI, even though the moderated-mediation index does not. The pattern is consistent with a reliably positive indirect effect across the observed familiarity range and an imprecisely estimated moderation slope. Study 3 brings the moderated-mediation test to conventional power. Figure 2 plots the *a*-path interaction directly.




Figure 2. Disclosure × consumer AI familiarity interaction on perceived option breadth (Study 2).

4.3. Discussion

Study 2 isolates the mechanism that explains the Study 1 hedonic-cell positive contrast. Disclosing the algorithmic origin of a hedonic curation lifts perceived option breadth, and that breadth inference in turn lifts purchase intention. The indirect effect is robust to bootstrapping and survives a head-to-head test against the two most plausible alternative mediators in the published literature, generalized trust and perceived competence. The opposite-sign direct effect ($c' = -0.12$ against a positive total effect $c = +0.20$) is a textbook inconsistent-mediation pattern [49], meaning the perceived-option-breadth mechanism more than absorbs the net disclosure effect on purchase intention. Study 2 also leaves three issues open: the moderated-mediation prediction by consumer AI familiarity is supported directionally but not at conventional statistical power; mentalizing was not included in the mediator race; and the single-factor option-breadth scale mixes search-side and display-side language at the item level. Study 3 addresses all three.

5. Cross-Study Summary Before Study 3

Before turning to Study 3, we summarize the cross-study pattern of the H1a hedonic-cell contrast across Studies 1 and 2 to establish what S3 is asked to extend. Table 7 places the focal results side by side. Study 1's hedonic cell yields an AI-versus-human purchase-intention difference of +0.64 ($d = +0.65$). Study 2's AI-versus-human pairwise contrast within the all-hedonic three-cell design yields +0.19 (Table 5 above). The two studies establish the hedonic-cell positive contrast in a single product category (music subscription) and across two independent samples. Study 3 is designed to extend the evidence on three fronts: a second hedonic category for stimulus generalizability, a bi-factor decomposition of the mediator, and a head-to-head mediator race that includes mentalizing. The moderated-mediation test of H3 is also brought to conventional statistical power.

Table 7. Cross-study H1a pattern (Studies 1 and 2).

Study	Stimulus	M_{AI}	M_{Human}	Diff.	Test Statistic	Cohen's d
Study 1 (hedonic cell)	Music subscription	5.47	4.83	+0.64	$t(112) = 3.46, p < 0.001$	+0.65
Study 2	Music subscription	5.22	5.03	+0.19	—	—

6. Study 3

6.1. Participants and Procedure

A sample of 600 adult online shoppers was recruited through Credamo, which was calibrated against a G*Power 3.1.9.7 [50] analysis identifying the Index of Moderated Mediation test as the binding power constraint at approximately 480–520 respondents ($\alpha = 0.05$, two-tailed, $1 - \beta = 0.80$, anchored on Study 2 $a = 0.55$, $b = 0.50$, and a medium moderator slope of 0.20). The 600-respondent recruitment target incorporated an approximately 15 percent buffer for exclusion-rule attrition. The a priori exclusion rules (attention-check failure, mismatched disclosure manipulation check, completion time outside the 60 to 1500 s window, and Voss, Spangenberg, and Grohmann [40] hedonic subscale mean below 4 for the assigned stimulus) retained an analytic sample of 519 respondents (52 percent female; $M_{age} = 31.8$ years), which was an 86.5 percent retention rate consistent with the corresponding rates in Studies 1 and 2. All inferential tests are two-tailed at $\alpha = 0.05$; the H1a hedonic-cell contrast (within and pooled across categories); the H2 search-side mediation, the H2a parallel-mediator bifurcation, and the H3 Index of Moderated Mediation [51] are treated as focal tests specified in advance. The study was approved under IRB protocol GNARIRB0013.

The design crossed curator disclosure (AI vs. human) with product type (Hedonic-A music subscription versus Hedonic-B video subscription) in a fully between-subjects allocation with Credamo's stratified random allocator targeting a cell size of approximately 130 (post-exclusion cell sizes were 132/127/135/125 across the four cells). The decision to hold the design within two hedonic categories rather than mix in a utilitarian cell reflects the role assigned to Study 3 in the cross-study program: Study 1 anchors the utilitarian-cell null contrast ($t(112) = -0.01, p = 0.99$), and Study 3 is allocated to deepening the hedonic-cell evidence through stimulus generalizability, bi-factor mediator measurement, mentalizing as a competing mediator, and the moderated-mediation repower. Consumer AI familiarity was measured before stimulus exposure to avoid post-treatment contamination, using the same four-item scale as in Study 2.

Participants gave informed consent, completed brief demographics, and completed the consumer-AI-familiarity and need-for-uniqueness measures pre-stimulus. Then, they were randomly assigned to one of the four stimulus screenshots. The participants viewed the assigned stimulus under a five-second minimum exposure timer enforced by the survey platform, completed a 30 s arithmetic distractor task, and then completed the disclosure-manipulation check, the Voss, Spangenberg, and Grohmann [40] hedonic-subscale manipulation check, and the substantive measurement battery. The substantive battery comprised seven scale blocks presented in fully randomized order across respondents—a procedural upgrade from the fixed-order presentation used in Studies 1 and 2. The blocks were search-side perceived option breadth (4 items), display-side perceived option breadth (4 items), mentalizing (3 items), trust in recommendation (4 items), AI competence (3 items), purchase intention (3 items), and attitude toward the recommendation (3 items). A single forced-choice behavioral item—"Would you click 'Add to Cart' for this recommendation now? (Yes/No)"—followed the substantive battery and serves as a complementary behavioral-intention dependent variable. Two attention checks [52] were embedded in the substantive battery. Participants were then debriefed and compensated through the Credamo platform.

6.2. Stimuli

Study 3 stimuli consist of four pixel-identical pairs of mock product-page screenshots, following the convention of Studies 1 and 2. Cells 1 and 2 present a music-subscription playlist landing page (Hedonic-A, re-used from Studies 1 and 2 for cross-study comparability); cells 3 and 4 present a drama-series recommendation surface on a generic video-subscription platform (Hedonic-B, newly developed for Study 3 to test stimulus generalizability). Within each Hedonic level, the AI and human cells differ only in a 14-character curator-disclosure badge ("AI-curated" or "Human-curated", rendered in simplified Chinese on the participant-facing stimulus) positioned directly above the recommendation block. All other visual elements—header, layout, recommendation items, ratings, calls-to-action, badge styling—are held constant within each Hedonic level. Neither stimulus identifies a real brand, displays subscription pricing, or includes social-proof cues such as user counts, isolating the disclosure cue as the single varying experimental element. Full stimulus exhibits appear in Appendix B.

6.3. Measures

The eight-item bi-factor perceived-option-breadth scale developed for Study 3 separates the search-side subdimension (four items targeting the consumer's inference about the size of the option pool the curator considered) from the display-side subdimension (four items targeting the consumer's perception of the visible recommendation surface). The full item set is in Appendix A. The three-item mentalizing scale was adapted from

Waytz, Heafner, and Epley [53], retaining the items targeting the agency dimension of mind perception (the dimension most relevant to a curator who is not embodied). All twelve new items (eight option-breadth and three mentalizing, plus the single behavioral-intention item) underwent the Brislin [54] two-translator forward-and-back translation procedure [55], and the final wording was reviewed by a three-expert content-validity panel comprising one consumer-psychology PhD, one IS/HCI PhD, and one quantitative-methods PhD [56]. The scale-level CVI/Avg exceeded the conventional 0.80 threshold for all three new subscales (search-side option breadth = 0.94; display-side option breadth = 0.92; mentalizing = 0.89); item-level CVI ≥ 0.67 [57] was achieved for all twelve items. Carry-over scales from Studies 1 and 2 (trust, AI competence, purchase intention, attitude, need for uniqueness, AI familiarity, and the Voss, Spangenberg, and Grohmann [40] hedonic subscale) retained their previously translated and validated wording. All scales used a seven-point Likert format with the same anchors as Studies 1 and 2.

6.4. Analysis Plan

The pre-registered analysis plan locked the following test specifications. H1a within each hedonic category is tested through an independent-samples *t*-test of purchase intention by disclosure. H1a pooled across categories is tested through an analysis of variance with disclosure and product type as factors. H2 is tested through PROCESS Model 4 [48] with disclosure as X, search-side option breadth as M, purchase intention as Y, and 10,000 percentile bootstrap resamples. H2a is tested through a parallel-mediator extension of the same PROCESS Model 4 with both search-side and display-side option breadth entered simultaneously. The mediator race tests search-side option breadth, display-side option breadth, mentalizing, trust, and AI competence as five parallel mediators in the same model. H3 is tested through PROCESS Model 7 with consumer AI familiarity (mean-centered) moderating the *a*-path of search-side option breadth. The behavioral dependent variable is analyzed through logistic regression of the Add-to-Cart indicator on disclosure, product type, search-side option breadth, and AI familiarity. The bi-factor structure of the eight-item option-breadth scale is tested through confirmatory factor analysis comparing three nested models: Model A (two correlated first-order factors), Model B (single general factor), and Model C (bi-factor model with a general factor plus two orthogonal specific factors). Model fit is evaluated against the Hu and Bentler [58] cutoffs (CFI ≥ 0.95 , RMSEA ≤ 0.08 , $\chi^2/df \leq 3.0$). All bootstrap and inferential procedures use a fixed random seed of 20260622 for reproducibility. Sensitivity analyses report the focal H1a pooled-hedonic and H2 indirect effects under three alternative inclusion regimes (analytic sample, full pre-exclusion sample, attention-pass-only intermediate sample), mirroring the transparency reporting in Study 1.

6.5. Measurement Validation

All scales cleared the conventional psychometric thresholds (Table 8). Cronbach's α ranged from 0.858 (Voss, Spangenberg, and Grohmann [40] hedonic subscale) to 0.962 (AI familiarity); composite reliability ranged from 0.904 to 0.972; average variance extracted ranged from 0.702 to 0.898. Heterotrait-monotrait ratios between any two of the substantive constructs ranged from 0.020 to 0.056, which is well below the conservative 0.85 threshold [45], supporting discriminant validity. The behavioral manipulation checks were also clean: 100 percent of retained respondents identified the assigned disclosure source correctly (a mechanical consequence of the exclusion rule), and the Voss, Spangenberg, and Grohmann [40] hedonic subscale returned means of 5.45 ($SD = 0.77$) for the music-subscription stimulus and 5.54 ($SD = 0.74$) for the video-

subscription stimulus, which were both comfortably above the 4.0 pass threshold and consistent across categories.

The confirmatory factor analysis on the eight-item option-breadth scale supported the bi-factor decomposition strongly. Model A (two correlated first-order factors) returned an excellent fit, $\chi^2 = 11.43$ on 19 df ($\chi^2/df = 0.60$), CFI = 1.002, TLI = 1.003, RMSEA = 0.000, GFI = 0.997, AGFI = 0.995. Model B (single general factor) returned a poor fit, $\chi^2 = 1636.27$ on 20 df ($\chi^2/df = 81.81$), CFI = 0.506, TLI = 0.308, RMSEA = 0.395. Model C (bi-factor with one general factor and two orthogonal specific factors) returned a fit essentially identical to Model A at the global indices, $\chi^2 = 7.10$ on 12 df ($\chi^2/df = 0.59$), CFI = 1.001, RMSEA = 0.000. The chi-square difference test between Models A and B was substantial, $\Delta\chi^2 = 1624.85$ ($\Delta df = 1$), $p < 0.0001$, clearly favoring the two-factor structure over the single-factor structure and providing direct empirical support for the bifurcation of perceived option breadth into search-side and display-side subdimensions. The fit of Models A and C is statistically indistinguishable at the global level; we use Model A's two correlated factors in the mediation analyses because the specific-factor scores in Model C are residualized on the general factor and would complicate interpretation of the indirect-effect pathways.

Table 8. Reliability and convergent/discriminant validity, Study 3 ($N = 519$).

Construct	k Items	α	CR	AVE
Search-side option breadth (OB-search)	4	0.929	0.950	0.826
Display-side option breadth (OB-display)	4	0.928	0.949	0.822
Mentalizing	3	0.898	0.936	0.831
Trust in recommendation	4	0.915	0.940	0.798
AI competence	3	0.884	0.928	0.812
Purchase intention	3	0.916	0.947	0.856
Attitude toward recommendation	3	0.891	0.932	0.821
Need for uniqueness	4	0.947	0.962	0.862
Consumer AI familiarity	4	0.962	0.972	0.898
Hedonic subscale [40]	4	0.858	0.904	0.702

6.6. Results

6.6.1. H1a Within and Pooled Across Hedonic Categories

In the music-subscription category, AI-disclosed recommendations produced higher purchase intention than human-disclosed recommendations, $M_{AI} = 5.21$ ($SD = 1.13$, $n = 132$); $M_{Human} = 4.65$ ($SD = 1.05$, $n = 127$); difference = +0.56; $t(257) = 4.11$, $p < 0.0001$; Cohen's $d = +0.51$. In the video-subscription category, the AI-versus-human contrast was also positive but smaller in magnitude, $M_{AI} = 5.04$ ($SD = 1.06$, $n = 135$); $M_{Human} = 4.73$ ($SD = 1.13$, $n = 125$); difference = +0.31; $t(258) = 2.29$, $p = 0.023$; Cohen's $d = +0.29$. Pooling across the two categories, AI-disclosed recommendations produced higher purchase intention than human-disclosed recommendations, $M_{AI} = 5.12$ ($SD = 1.10$, $n = 267$); $M_{Human} = 4.69$ ($SD = 1.09$, $n = 252$); difference = +0.43; $t(517) = 4.52$, $p < 0.0001$; Cohen's $d = +0.40$. The 2×2 ANOVA on the full Study 3 sample yielded a significant main effect of disclosure, $F(1, 515) = 20.46$, $p < 0.0001$, partial $\eta^2 = 0.038$, with a non-significant main effect of product type ($F = 0.30$, $p = 0.58$) and non-significant Disclosure $\times$ Product type interaction ($F = 1.65$, $p = 0.20$). The non-significant interaction term indicates that the AI-versus-human lift does not differ statistically between the two hedonic categories, supporting the category-internal stability of the H1a effect even though the music-category point estimate is somewhat larger than the video-category point estimate. The cross-category stability provides direct evidence on the single-stimulus concern: the H1a hedonic-cell contrast survives in a second hedonic product category. Sensitivity analyses on the full pre-exclusion sample ($N = 600$) and the attention-pass-

only intermediate sample ($N = 540$) returned essentially identical pooled-hedonic effects, $d = +0.41$ and $d = +0.40$, respectively, mirroring the robustness pattern reported for Study 1.

6.6.2. H2 Single-Mediator and H2a Parallel-Mediator Decomposition

The H2 single-mediator test with search-side perceived option breadth as the sole mediator returned an indirect effect of $ab = +0.297$, 95% percentile bootstrap CI $[+0.206, +0.400]$, excluding zero. The result reproduces the Study 2 mediation pathway with a meaningfully tighter confidence interval (Study 2 $ab = +0.32$, 95% CI $[+0.16, +0.49]$; Study 3 CI width is approximately 39 percent narrower). The H2a parallel-mediator test, entering both search-side and display-side perceived option breadth simultaneously, returned indirect effects of $ab_{\text{search}} = +0.308$, 95% CI $[+0.217, +0.412]$ and $ab_{\text{display}} = +0.120$, 95% CI $[+0.065, +0.188]$. Both subdimensions carry a positive indirect effect, but the search-side subdimension's indirect effect is approximately 2.6 times the magnitude of the display-side subdimension's, and the two confidence intervals do not overlap. The direct effect after controlling for both mediators is near zero, $c' = +0.006$, 95% CI $[-0.172, +0.183]$. H2a is supported: the search-side subdimension is the dominant mediator; the display-side subdimension contributes a smaller but reliably positive secondary pathway.

6.6.3. Mediator Race Against Mentalizing, Trust, and Competence

The five-mediator race tested search-side option breadth, display-side option breadth, mentalizing, trust in recommendation, and AI competence as parallel mediators of the AI-disclosure effect on purchase intention. Search-side option breadth carried the largest indirect effect, $ab = +0.312$, 95% CI $[+0.222, +0.415]$, excluding zero. Display-side option breadth carried a smaller positive indirect effect, $ab = +0.115$, 95% CI $[+0.063, +0.177]$, also excluding zero. Mentalizing carried a small *negative* indirect effect, $ab = -0.078$, 95% CI $[-0.136, -0.025]$, excluding zero on the negative side: the AI-disclosed cell elicited lower mentalizing inference than the human-disclosed cell (a -path $= -0.26$), and mentalizing positively predicted purchase intention (b -path $= +0.30$), so the disclosure $\rightarrow$ mentalizing $\rightarrow$ PI pathway is negative on net. The pattern is consistent with the direction Wien and Peluso [6] identify: human recommenders elicit stronger mentalizing inference than AI recommenders. The substantive interpretation is that the mentalizing pathway operates in our data in the direction Wien and Peluso [6] identify, but it is dominated in net effect by the search-side option-breadth pathway that does most of the explanatory work. Trust in recommendation ($ab = +0.003$, 95% CI $[-0.021, +0.030]$) and AI competence ($ab = +0.015$, 95% CI $[-0.009, +0.046]$) carried no detectable indirect effect, reproducing the Study 2 mediator-race result. The total indirect effect across all five mediators was $+0.368$, 95% CI $[+0.241, +0.502]$. Table 9 summarizes the five paths.

Table 9. Mediator race results, Study 3 (X = AI vs. Human, M = five candidate mediators in parallel, Y = purchase intention).

Mediator	<i>a</i> -Path	<i>b</i> -Path	Indirect <i>ab</i>	95% Bootstrap CI	CI excl. 0
Search-side option breadth	+0.71	+0.44	+0.312	$[+0.222, +0.415]$	✓
Display-side option breadth	+0.40	+0.29	+0.115	$[+0.063, +0.177]$	✓
Mentalizing	−0.26	+0.30	−0.078	$[-0.136, -0.025]$	✓ (neg.)
Trust in recommendation	+0.02	+0.13	+0.003	$[-0.021, +0.030]$	—
AI competence	+0.11	+0.14	+0.015	$[-0.009, +0.046]$	—
Direct effect c'	—	—	+0.066	$[-0.104, +0.235]$	—
Total indirect	—	—	+0.368	$[+0.241, +0.502]$	✓

Note: ✓ indicates that the 95% percentile bootstrap confidence interval excludes zero; ✓ (neg.) indicates a reliably negative indirect effect. Bootstrap resamples = 10,000; seed = 20260622.

6.6.4. H3 Moderated Mediation by Consumer AI Familiarity

The H3 test through PROCESS Model 7 with consumer AI familiarity moderating the *a*-path of search-side option breadth returned an index of moderated mediation of +0.065, 95% percentile bootstrap CI [+0.014, +0.118], excluding zero. The conditional indirect effects rose monotonically across the consumer-AI-familiarity distribution: +0.205, 95% CI [+0.099, +0.324] at -1 SD; +0.297, 95% CI [+0.206, +0.400] at the mean; and +0.389, 95% CI [+0.268, +0.526] at $+1$ SD. Each conditional indirect effect excludes zero on its own confidence interval, and the moderation slope itself also excludes zero. H3 is supported: the search-side mediation is amplified at higher levels of consumer AI familiarity. The direction and magnitude are consistent with the directional pattern observed in Study 2; the difference between studies is that the moderated-mediation index now reliably excludes zero at the Study 3 sample size, which the a priori power calculation calibrated for this test. Figure 3 plots the conditional indirect effect across the observed range of consumer AI familiarity.




Figure 3. Conditional indirect effect of AI-disclosure on purchase intention via search-side option breadth, across consumer AI familiarity (Study 3). Shaded bands are 95% percentile bootstrap CIs.

6.6.5. Behavioral Dependent Variable

The logistic regression of the Add-to-Cart binary indicator on disclosure, product type, the Disclosure $\times$ Product type interaction, search-side option breadth, and consumer AI familiarity ($N = 519$, pseudo $R^2 = 0.049$) returned the following pattern. Disclosure had a non-significant direct effect on Add-to-Cart once search-side option breadth was in the model ($\beta = +0.31$, $OR = 1.36$, $p = 0.28$). Search-side option breadth was a strong predictor of Add-to-Cart ($\beta = +0.42$, $OR = 1.52$, $p < 0.0001$). Product type, the Disclosure $\times$ Product type interaction, and consumer AI familiarity were non-significant predictors at conventional thresholds. The pattern is consistent with the mediation framework: the AI-disclosure cue does not predict Add-to-Cart directly once the mediator is in the model; the mediator does. Cell-level Add-to-Cart rates were 75.0 percent (AI $\times$ music), 71.9 percent (AI $\times$ video), 60.6 percent (Human $\times$ music), and 63.2 percent (Human $\times$ video), which is consistent with the cell-level purchase-intention pattern reported above.

6.7. Discussion

Study 3 deepens the cross-study evidence on four fronts. The hedonic-cell AI-versus-human positive contrast reproduces in a second product category, and the pooled hedonic-cell effect ($d = +0.40$) sits within a tighter confidence interval than Study 1's single-category estimate. The bi-factor decomposition of perceived option breadth into search-side and display-side subdimensions is empirically supported by confirmatory factor analysis, and the parallel-mediator mediation analysis identifies the search-side subdimension as the dominant indirect pathway with a smaller positive contribution from the display-side subdimension. Mentalizing operates in the direction Wien and Peluso [6] document—the AI-disclosed cell elicits lower mentalizing inference than the human-disclosed cell, and mentalizing positively predicts purchase intention—but it is dominated in net effect by the search-side option-breadth pathway. The moderated-mediation test of H3 clears the conventional significance threshold at Study 3's sample size, supporting the prediction that consumer AI familiarity amplifies the search-side indirect effect. A single forced-choice behavioral-intention item complements the self-reported purchase-intention measure, and the logistic regression of Add-to-Cart on the mediator and the disclosure cue is consistent with the mediation framework. Section 7 develops the cross-study theoretical, regulatory, and managerial implications.

7. General Discussion

Across three online experiments, AI-disclosed recommendations outperform human-disclosed recommendations on purchase intention in hedonic e-commerce, while no analogous lift emerges for utilitarian products. The cross-study effect is mechanism-attributable through search-side perceived option breadth: a consumer inference about the size of the option pool the curator considered upstream, which was conceptually and empirically distinct from the social-evaluation mediators (competence, warmth, generalized trust, anthropomorphism, mentalizing) that have so far dominated the AI-disclosure literature. Within the algorithmic-symbiosis paradigm articulated by the editors of this topic [59,60], the findings re-cast AI-curation disclosure from a trust liability to be minimized into a strategic communication lever that interacts with the cognitive structure of the consumption category.

7.1. Theoretical Contributions

The first contribution is a falsifiable boundary condition on the algorithm-aversion-in-subjective-domains pattern that has anchored consumer–AI research for the past decade. Castelo, Bos, and Lehmann [3] anchored the prediction that subjective tasks invite algorithm aversion; Wien and Peluso [6] extended the prediction to retail recommendation; and Qin et al.'s (2025) [16] capability–personalization meta-analysis reaches a similar directional conclusion for high-personalization contexts. The present cross-study evidence stands directly opposite this convergent prediction within the hedonic-recommendation case. The qualification is conditional rather than universal. The hedonic-cell AI-versus-human positive contrast survives in two studies (Study 1 and Study 3) and two product categories (music subscription and video subscription); the utilitarian-cell null contrast reproduces (Study 1); and the underlying mechanism (search-side perceived option breadth operating through preference-formative consumption) generates a falsifiable empirical signature—the Disclosure $\times$ Product-type interaction—rather than a brute reversal of the aversion direction.

The second contribution is a mediator that is conceptually and empirically distinct from competence, warmth, generalized trust, anthropomorphism, and mentalizing. Search-side perceived option breadth is an attribution about the curator's *computational input pool*—not about the curator as a social agent. The Study 2 mediator race establishes the

distinction against trust and competence; the Study 3 mediator race extends the distinction against mentalizing. The mentalizing pathway operates in our data in the direction Wien and Peluso [6] identified, but it is dominated in net effect by the search-side option-breadth pathway that does most of the explanatory work in our hedonic categories. We read this as the search-side pathway *augmenting* rather than *replacing* the mentalizing pathway in the AI-disclosure literature with the relative weight of the two pathways being an empirical question we have begun to answer for hedonic e-commerce specifically. Future research can compare the two pathways' relative dominance in adjacent contexts—service recommendation, content creation, expert advice—where the cognitive structure of the consumption decision may shift the balance.

The third contribution is a theoretical bridge between two types of literature that have, until now, developed in parallel. Algorithm-aversion research has treated the AI source primarily through a social-evaluation lens, while assortment-perception research has treated curator effort and variety inference as central drivers of consumer satisfaction [10,11]. Bringing the second literature to bear on the first opens an empirical space in which the AI-disclosure effect can be positive in domains where the social-agent mediators would predict it to be negative. The bridging construct—perceived option breadth, decomposed into search-side and display-side subdimensions in our refined eight-item bi-factor scale—is one such instrument. The preference-formative-versus-confirming distinction we develop in Section 2.2 supplies the deeper theoretical reason that breadth is diagnostic in some categories and not in others, and it connects the AI-disclosure question to the broader preference-construction tradition in consumer behavior [15,25].

7.2. Regulatory Implications

The disclosure-research orthodoxy assumes that transparency creates friction: when consumers learn that a message has been algorithmically produced or curated, persuasion-knowledge defences activate and consumer responses dampen [38,61]. Mandatory algorithm-disclosure regulation—Cyberspace Administration of China's (2022) [7] Algorithm Recommendation Provisions, Cyberspace Administration of China's (2023) [8] Interim Measures for Generative AI Services, and the European Union [9] Article 50 transparency obligation—is built on this orthodoxy: regulators reason that consumers, once informed of algorithmic mediation, will exercise more critical scrutiny. The present findings complicate this assumption in a non-trivial way. In hedonic e-commerce categories, the same disclosure cue the friction account predicts to dampen persuasion instead lifts purchase intention through a parallel inferential route (search-side option breadth) that the friction account does not consider. The implication is that a disclosure regime designed to attenuate algorithmic persuasion in some categories may simultaneously amplify it in others in a way that the categorical reasoning of one-size-fits-all transparency rules does not anticipate. The right response is not to weaken disclosure but to recognize that what the consumer infers from a disclosure cue is category-dependent, and that uniform transparency rules carry under-theorized interaction effects with the cognitive structure of the consumption category. A more nuanced regulatory design would distinguish between categories in which disclosure dampens persuasion (a feature) and categories in which it amplifies persuasion (also a feature, when the underlying algorithmic curation genuinely delivers the breadth the disclosure signals). We frame this as a theoretical prediction derived from three online experiments rather than a recommendation that any particular regulator should alter any particular rule; field replication outside the categories we tested and outside the Chinese consumer panel context is required before policy translation.

7.3. Managerial Implications

For e-commerce platforms operating hedonic surfaces—music subscriptions, fashion feeds, video subscriptions, content recommendations, gift discovery, experiential-goods catalogues—the present evidence supports disclosing AI curation rather than concealing it, pending field replication on consequential consumer behavior. A simple curation-label badge of the kind tested here (“AI-curated for you”) produced a medium-to-large lift in self-reported purchase intention in our hedonic cells with no analogous effect in the utilitarian cell. For platforms operating at scale, the disclosure-badge intervention is a user-experience switch with essentially no marginal cost, and the present mechanism evidence implies that pairing the badge with an explicit breadth cue (language that makes the algorithm’s search scope salient to the consumer) should amplify rather than redundantly duplicate the effect. The stimulus copy used in our Study 2 includes the breadth pairing as a natural extension; field deployments combining the disclosure badge with an explicit breadth statistic should expect compound rather than competing effects, conditional on the platform’s underlying algorithmic curation genuinely delivering the breadth the disclosure signals. For products in which consumers arrive with a tight technical specification (printers, batteries, replacement parts, low-differentiation utilitarian categories), the disclosure cue produced no detectable lift in our data; allocating user-experience real estate to a prominent curation badge in these categories is unlikely to repay the attention cost. The decision rule is therefore directional rather than universal: on hedonic surfaces, disclose AI curation and accompany the disclosure with a breadth cue, conditional on the underlying algorithmic infrastructure being able to back the signal; on utilitarian surfaces, treat the disclosure decision as neutral and use the available real estate for specification-relevant information. We treat each of these recommendations as conditional on field replication outside the Chinese consumer panel context.

7.4. Limitations and Future Research

Several limitations bound the present findings. The first concerns stimulus generality: our hedonic stimuli cover two product categories (music subscription and video subscription) and the utilitarian-cell evidence comes from a single product category (wireless mouse). The two-category coverage in the hedonic cell is an improvement over single-stimulus designs but does not exhaust the hedonic category space (fashion, food, travel content, fragrance, gift discovery, experiential goods remain to be tested). A meta-analytic synthesis of disclosure $\times$ product-type interactions across categories, or a multi-category replication in a single study, is the natural medium-term project.

The second limitation concerns the dependent variable. The primary dependent variable in all three studies is self-reported purchase intention; Study 3 supplements with a single forced-choice behavioral-intention item (“Add to Cart”). Field experiments—A/B tests of the curation-label badge on partner e-commerce platforms with consequential transactional outcomes (click-through, conversion, basket value)—would translate the present evidence into the revenue effects managers ultimately care about. The breadth-cue managerial extension proposed in the preceding subsection could be tested concurrently in such a field design.

The third limitation concerns the generality of our sample. All three experiments draw from a single Chinese consumer panel; cross-cultural replication is needed before the effect can be claimed as general. Qin et al. [16] observe in their meta-analysis that AI aversion is more pronounced in countries with high education and internet penetration, which suggests that the present effect’s strength may co-vary with consumer AI familiarity at the population level. A Western-sample replication would be informative both as a generality test and as a calibration point on the cross-cultural picture. The H3 moderated-mediation

pattern offers a partial within-sample test of this generality conjecture—the indirect effect is stronger for high-AI-familiarity respondents—but the within-sample evidence cannot substitute for a cross-cultural replication.

The fourth limitation concerns common-method variance. Both the mediator (search-side perceived option breadth) and the primary outcome (purchase intention) were measured within the same survey wave through self-report. The Study 1 H1a hedonic-cell finding does not rely on the option-breadth measure (it is an ANOVA on purchase intention alone), which provides some protection against common-method-variance interpretation. The Study 2 and Study 3 mediation tests should, however, be regarded as preliminary until a separated-wave design or an objective behavioral outcome replicates the indirect effect [62].

The fifth limitation concerns the scope of the mentalizing competing-mediator test. Study 3 fields a three-item agency-focused short form adapted from Waytz, Heafner, and Epley [53], retaining the items most relevant to a non-embodied recommendation system. A longer or alternative mentalizing scale [63,64] could yield a different estimate of the mentalizing pathway's magnitude. The cross-study direction of the mentalizing effect we observe is consistent with Wien and Peluso [6]; the magnitude question is one that future work with a richer mentalizing instrument can address.

Finally, we deliberately did not vary the generative aspect of AI. Synthetic-content disclosures of the kind raised by De Bruyn, Viswanathan, Beh, Brock, and Von Wangenheim [65] and De Freitas, Nave, and Puntoni [66] likely activate concerns distinct from the recommender-AI case we study here, including hallucination, originality, and authenticity. The breadth mechanism we identify is, on our reading, a recommender-AI mechanism rather than a generative-AI mechanism, and the boundary between the two is itself an empirical question worth pursuing.

8. Conclusions

For hedonic e-commerce categories, disclosing that recommendations are AI-curated is not a trust liability to be minimized. It is a strategic communication lever that raises purchase intention through search-side perceived option breadth—a mechanism distinct from the social-evaluation mediators that have dominated the AI-disclosure literature and operationalized through an empirical decomposition of perceived option breadth into search-side and display-side subdimensions in a refined eight-item bi-factor scale. The finding constitutes a falsifiable boundary condition on the algorithm-aversion-in-subjective-domains pattern established by a decade of prior work, supplies a novel mechanism that bridges algorithm-aversion research with the assortment-perception tradition, and qualifies the friction-creates-attenuation assumption that anchors much of the contemporary algorithmic-disclosure regulatory architecture. Practitioners deploying AI recommenders on hedonic surfaces should disclose AI curation and pair the disclosure with explicit breadth cues conditional on the underlying algorithmic infrastructure genuinely delivering the breadth the disclosure signals; field evidence is the natural next step.

Author Contributions: Conceptualization, Y.W., T.M., Y.D. and L.Z.; methodology, T.M.; software, Y.W. and T.M.; writing—original draft preparation, T.M., Y.D. and Y.W.; writing—review and editing, T.M., Y.W., Y.D. and L.Z.; project administration, L.Z.; funding acquisition, L.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Guizhou Provincial Basic Research Program (Natural Science, No. [2025]237).

Institutional Review Board Statement: All three studies were approved by the Ethics Committee of Guian New Area Institute, Guizhou University of Finance and Economics (Approval Code GNARIRB0013, 15 October 2025).

Informed Consent Statement: Online informed consent was obtained from all participants involved in the study prior to participation.

Data Availability Statement: De-identified analytic data supporting the conclusions of this article will be deposited on the Open Science Framework upon acceptance.

Acknowledgments: Authors contributed equally to this work and share first authorship.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Scale Items

The scales below were administered in simplified Chinese on a seven-point Likert response format (1 = “strongly disagree,” 7 = “strongly agree”) unless otherwise noted. English wording is reproduced for the international reader; back-translation by an independent bilingual research assistant under the Brislin [54] procedure confirmed semantic equivalence between the English and Chinese versions for all newly developed items. Carry-over scales from Studies 1 and 2 retain the previously validated Chinese wording.

Appendix A.1. Search-Side Perceived Option Breadth (4 Items; Study 3; $\alpha = 0.929$)

The items refer to the curator’s *upstream search effort* rather than the displayed recommendation surface. Items developed for this paper as the search-side subdimension of the bi-factor option-breadth scale.

1. The recommendation system considered a very large number of products before showing me these recommendations.
2. I believe the system searched through an extensive set of options when generating these recommendations.
3. Before making these recommendations, the system likely evaluated many more products than what I see here.
4. These recommendations reflect a wide search across the available product space.

Appendix A.2. Display-Side Perceived Option Breadth (4 Items; Study 3 $\alpha = 0.928$)

The items refer to the visible recommendation surface rather than the upstream considered pool. Items developed for this paper as the display-side subdimension of the bi-factor option-breadth scale.

1. The recommendations shown to me cover a wide range of options.
2. The visible set of recommendations is varied.
3. What I see here spans many different choices.
4. The displayed recommendations represent a broad selection.

Appendix A.3. Single-Factor Perceived Option Breadth (4 Items; Studies 1 and 2)

Items developed for this paper for Studies 1 and 2. The four items target the search-effort attribution underlying perceived option breadth but mix search-side and display-side language at the item level; the Study 3 bi-factor scale (Appendices A.1 and A.2) refines the measurement.

1. The recommendations consider a wide range of options.
2. Many alternatives were evaluated before showing these recommendations.
3. The selection draws from a broad pool of products.
4. These recommendations were selected from a large variety of items.

Appendix A.4. Mentalizing (3 Items; Study 3 $\alpha = 0.898$)

Adapted from Waytz, Heafner, and Epley [53], retaining the agency-dimension items most relevant to a non-embodied recommendation system.

1. The recommendation system seemed to have its own way of thinking when choosing what to show me.
2. The recommendation system appeared to make intentional choices about what to recommend.
3. The recommendation system seemed to know what it was doing.

Appendix A.5. Purchase Intention (Dodds, Monroe, and Grewal [41]; 3 Items; $\alpha = 0.885/0.85/0.916$)

1. I would consider buying one of these items.
2. The likelihood of me purchasing one of these items is high.
3. My willingness to buy one of these items is high.

Appendix A.6. Trust in Recommendation (Komiak and Benbasat [34]; 4 Items; $\alpha = 0.905/0.91/0.915$)

The original Komiak and Benbasat scale contains seventeen items across three dimensions; we selected four items most directly applicable to a recommendation-page context and reworded item 4 from “competent” to “credible” to avoid lexical overlap with the AI-competence construct also measured in our instrument.

1. I trust the recommendations on this page.
2. The recommendations seem reliable.
3. I would feel safe acting on these recommendations.
4. I believe the recommendations are credible.

Appendix A.7. Consumer AI Familiarity (Adapted from Logg, Minson, and Moore [21]; 4 Items; $\alpha = 0.950/0.95/0.962$)

1. I am familiar with how AI tools work.
2. I use AI tools (chatbots, recommendation engines) regularly.
3. I understand how AI generates personalised recommendations.
4. AI is part of my everyday digital life.

Appendix A.8. Perceived AI Competence (Adapted from Fiske, Cuddy, and Glick [31]; 3 Items; $\alpha = 0.876/0.86/0.884$)

1. The curator behind the recommendations is competent.
2. The curator behind the recommendations is skillful.
3. The curator behind the recommendations is intelligent.

Appendix A.9. Attitude Toward Recommendation (MacKenzie and Lutz [42]; 3 Items, 7-Point Semantic Differential; Study 3 $\alpha = 0.891$)

- Bad/good
- Unfavorable/favourable
- Negative/positive

Appendix A.10. Hedonic/Utilitarian Product Perception (Voss, Spangenberg, and Grohmann [40]; 7-Point Semantic Differential; Studies 1 and 3)

- Hedonic: not fun/fun · not exciting/exciting · dull/thrilling · not delightful/delightful · not enjoyable/enjoyable

- Utilitarian: not useful/useful · impractical/practical · not functional/functional · unnecessary/necessary · ineffective/effective

Appendix A.11. Need for Uniqueness (Tian, Bearden, and Hunter [43]; Short Form; Study 1 Only; $\alpha = 0.852$; Item 4 Reverse-Coded)

1. I often think of things that please me but are different from what others would buy.
2. I tend to seek products that express my personal individuality.
3. I am attracted to rare items that few others have.
4. I avoid products that are popular with most people. (R)

Appendix A.12. Behavioral Click-Through Intention (Single Item, Study 3)

“Would you click ‘Add to Cart’ for this recommendation now? (Yes/No)”

Coded 1 = Yes, 0 = No for logistic-regression analysis.

Appendix A.13. Curator-Disclosure Manipulation Check (Forced-Choice)

“Who curated the recommendations you just saw?” Options: “AI/Editors/Not stated/Not sure.”

Appendix A.14. Attention Checks (Oppenheimer, Meyvis, and Davidenko [52] IMC Pattern; Two Items)

1. “Please select ‘Strongly disagree’ for this item.”
2. “For quality-control purposes, please select ‘Somewhat agree’ here.”

Appendix B. Stimulus Materials

Appendix B.1. Studies 1 and 2 Stimuli (Four Cells of the Two-by-Two Design Plus Study 2 No-Label Baseline)

Each respondent saw one product-page screenshot embedded inline in the survey. The four cells of the two-by-two design, together with the Study 2 no-label baseline, were constructed as custom-built mock-up product pages in simplified Chinese. The cells share an identical visual chassis (navigation bar, breadcrumb, hero region, related-product grid, pricing, and call-to-action), and they differ only in the curation-label badge immediately above the hero title. The AI-disclosed cells display a violet-gradient badge reading “AI-curated for you” (rendered in simplified Chinese in the stimulus); the human-disclosed cells display an amber-gradient badge reading “Curated by our editors” (rendered in simplified Chinese); and the Study 2 no-label baseline removes the badge entirely while leaving every other element pixel-identical. The hedonic stimulus pairs render a music-subscription playlist landing page. The hero playlist is presented with album art that depicts a moonlit-sky vinyl record, and the related-product grid below shows six visually distinct mood-themed playlists. The utilitarian stimulus pairs render a wireless-mouse product page. The hero product is presented with a detailed product render and a specification list (2.4 GHz wireless connection, 18-month battery life, five DPI levels, multi-platform compatibility, 12-month warranty), and the related-product grid below shows six other wireless-mouse options.




Figure A1. Side-by-side comparison of the four experimental cells in Studies 1 and 2. The two left panels (A,C) display the AI-disclosed badge; the two right panels (B,D) display the human-disclosed badge. Panels (A,B) show the hedonic stimulus (music-subscription playlist); panels (C,D) show the utilitarian stimulus (wireless mouse). All other visual elements are pixel-identical within each row. The stimulus interface text visible in each panel (e.g., product titles, playlist names, category labels, and call-to-action buttons) was rendered in simplified Chinese as presented to Chinese-panel participants; English glosses of the badge wording ("AI-curated for you" vs. "Curated by our editors") and stimulus content are provided in the body text and in Appendix B.

Appendix B.2. Study 3 Stimuli (Four Cells of the Disclosure $\times$ Hedonic-Category Design)

Study 3 stimuli comprise four pixel-identical pairs covering the music-subscription playlist landing page (carried over from Studies 1 and 2 as Hedonic-A) and a drama-series recommendation surface on a generic video-subscription platform (Hedonic-B, newly developed). Within each Hedonic level, the AI and human cells differ only in a 14-character curator-disclosure badge positioned directly above the recommendation block. The video-subscription stimulus presents a 2×2 grid of four drama-series posters above the fold, which are each accompanied by a brief tagline, episode count, and star rating; the music-subscription stimulus retains its five-track playlist hero layout. Neither stimulus identifies a real brand, displays subscription pricing, or includes social-proof cues such as user counts. Scenario primer text presented to participants before stimulus exposure positions both categories as leisure consumption—end-of-day music listening for the music subscription, weekend-evening drama-series viewing for the video subscription—without mentioning price, technical features, or any specification.




Figure A2. Side-by-side comparison of the four Study 3 cells. The two top panels show the music-subscription stimulus (Hedonic-A) under the AI-curated and human-curated badges; the two bottom panels show the video-subscription stimulus (Hedonic-B) under the same two badges. Each panel's stimulus interface text (product titles, playlist and drama-series names, navigation labels, and call-to-action buttons) was rendered in simplified Chinese as presented to Chinese-panel participants; English glosses of the badge wording ("AI-curated" vs. "Human-curated") and the cell labels—Cell 1: AI × Music (Hedonic-A); Cell 2: Human × Music (Hedonic-A); Cell 3: AI × Video (Hedonic-B); Cell 4: Human × Video (Hedonic-B)—are shown adjacent to each panel.

A 30-respondent pretest confirmed that each AI/human badge produced a 90-percent-or-higher correct identification rate on the forced-choice manipulation check, that the hedonic stimuli scored higher on the Voss, Spangenberg, and Grohmann [40] hedonic semantic differential than on the utilitarian semantic differential (and vice versa for the utilitarian stimulus in Studies 1 and 2), and that the perceived-option-breadth items were comprehensible to non-expert readers.

Appendix C. Construct Operational Definitions

The table below summarizes the conceptual and operational definitions of every measured construct across the three studies, the source of each scale, and the construct's status across the cross-study program.

Table A1. Operational definitions of measured constructs across the three studies.

Construct	Conceptual definition	Operational definition	Source	Items	Status
OB-search	Consumer's inference about the size of the option pool the curator considered, attributed to the curator's search effort, distinct from displayed variety	Four 7-point Likert items referencing the upstream curator search	Developed for this paper (S3 bi-factor scale)	4 (A.1)	Refined for S3
OB-display	Consumer's perception of the breadth, variety, and category coverage of the visible recommendation surface	Four 7-point Likert items referencing the visible surface	Developed for this paper (S3 bi-factor scale)	4 (A.2)	New for S3
OB (single-factor)	Pre-bifurcation predecessor mixing search-side and display-side language at the item level	Four 7-point Likert items	Developed for this paper (Studies 1, 2)	4 (A.3)	S1, S2 only
Mentalizing	Attribution of mental states (agency dimension) to the curator agent	Three 7-point Likert items focused on the agency dimension	Adapted from Waytz, Heafner, and Epley [53]	3 (A.4)	New for S3
Trust in recommendation	Consumer's confidence in the recommendation as a reliable basis for the decision	Four 7-point Likert items	Komiak and Benbasat [34]; item 4 reworded "competent" → "credible"	4 (A.6)	Anchored
Perceived AI competence	Consumer's inference about the AI platform's general technical competence	Three 7-point Likert items	Adapted from Fiske, Cuddy, and Glick [31]	3 (A.8)	Anchored
Consumer AI familiarity	Consumer's prior exposure to and self-rated knowledge of AI-enabled products	Four 7-point Likert items; measured pre-stimulus in S2 and S3	Adapted from Logg, Minson, and Moore [21]	4 (A.7)	Anchored
Purchase intention	Consumer's stated likelihood of purchasing the recommended product	Three 7-point Likert items	Dodds, Monroe, and Grewal [41]	3 (A.5)	Anchored
Attitude toward recommendation	Consumer's overall affective evaluation of the recommendation	Three 7-point semantic-differential items	MacKenzie and Lutz [42]	3 (A.9)	Anchored
Need for uniqueness	Dispositional tendency to seek differentiation through consumer choices	Four 7-point Likert items (short form)	Tian, Bearden, and Hunter [43]	4 (A.11)	Study 1 only
Hedonic/Utilitarian product perception	Consumer's perception of the product along the hedonic and utilitarian dimensions	7-point semantic-differential items	Voss, Spangenberg, and Grohmann [40]	5 + 5 (A.10)	S1 + S3 (HED in S3)
Behavioral click-through intention	Binary indicator of whether the consumer would proceed to add the recommended product to the cart	Single forced-choice item	Developed for this paper	1 (A.12)	New for S3
Curator-disclosure manipulation check	Cell-correct identification of the disclosed curator source	Forced-choice item	Developed for this paper	1 (A.13)	Anchored
Attention checks	Cell-irrelevant items designed to detect satisficing	Two embedded items	Oppenheimer, Meyvis, and Davidenko [52] IMC pattern	2 (A.14)	Anchored

References

- Davenport, T.; Guha, A.; Grewal, D.; Bressgott, T. How artificial intelligence will change the future of marketing. *J. Acad. Mark. Sci.* **2020**, *48*, 24–42. [\[CrossRef\]](#)
- Huang, M.-H.; Rust, R.T. The caring machine: Feeling AI for customer care. *J. Mark.* **2024**, *88*, 1–23. [\[CrossRef\]](#)
- Castelo, N.; Bos, M.W.; Lehmann, D.R. Task-dependent algorithm aversion. *J. Mark. Res.* **2019**, *56*, 809–825. [\[CrossRef\]](#)
- Dietvorst, B.J.; Simmons, J.P.; Massey, C. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *J. Exp. Psychol. Gen.* **2015**, *144*, 114–126. [\[CrossRef\]](#) [\[PubMed\]](#)
- Longoni, C.; Bonezzi, A.; Morewedge, C.K. Resistance to medical artificial intelligence. *J. Consum. Res.* **2019**, *46*, 629–650. [\[CrossRef\]](#)
- Wien, A.H.; Peluso, A.M. Influence of human versus AI recommenders: The roles of product type and cognitive processes. *J. Bus. Res.* **2021**, *137*, 13–27. [\[CrossRef\]](#)
- Cyberspace Administration of China. *Provisions on the Administration of Algorithm Recommendation in Internet Information Services*; Cyberspace Administration of China: Beijing, China, 2022.
- Cyberspace Administration of China. *Interim Measures for the Management of Generative Artificial Intelligence Services*; Cyberspace Administration of China: Beijing, China, 2023.
- European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). In *The Official Journal of the European Union*; European Union: Brussels, Belgium, 2024; L 1689.
- Diehl, K.; Poynor, C. Great expectations?! Assortment size, expectations, and satisfaction. *J. Mark. Res.* **2010**, *47*, 312–322. [\[CrossRef\]](#)
- Kahn, B.E.; Wansink, B. The influence of assortment structure on perceived variety and consumption quantities. *J. Consum. Res.* **2004**, *30*, 519–533. [\[CrossRef\]](#) [\[PubMed\]](#)
- Botti, S.; Iyengar, S.S. The dark side of choice: When choice impairs social welfare. *J. Public Policy Mark.* **2006**, *25*, 24–38. [\[CrossRef\]](#)
- Hoch, S.J. Product experience is seductive. *J. Consum. Res.* **2002**, *29*, 448–454. [\[CrossRef\]](#) [\[PubMed\]](#)
- Khan, U.; Dhar, R.; Wertenbroch, K. A behavioral decision theory perspective on hedonic and utilitarian choices. In *Inside Consumption: Consumer Motives, Goals, and Desires*; Ratneshwar, S., Mick, D.G., Eds.; Routledge: Oxfordshire, UK, 2005; pp. 144–165.
- Slovic, P. The construction of preference. *Am. Psychol.* **1995**, *50*, 364–371. [\[CrossRef\]](#)
- Qin, X.; Zhou, X.; Chen, C.; Wu, D.; Zhou, H.; Dong, X.; Cao, L.; Lu, J.G. AI aversion or appreciation? A capability–personalization framework and a meta-analytic review. *Psychol. Bull.* **2025**, *151*, 580–599. [\[CrossRef\]](#) [\[PubMed\]](#)
- Prahl, A.; Van Swol, L. Understanding algorithm aversion: When is advice from automation discounted? *J. Forecast.* **2017**, *36*, 691–702. [\[CrossRef\]](#)
- Yeomans, M.; Shah, A.; Mullainathan, S.; Kleinberg, J. Making sense of recommendations. *J. Behav. Decis. Mak.* **2019**, *32*, 403–414. [\[CrossRef\]](#)
- Dietvorst, B.J.; Simmons, J.P.; Massey, C. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Manag. Sci.* **2018**, *64*, 1155–1170. [\[CrossRef\]](#)
- Jussupow, E.; Benbasat, I.; Heinzl, A. Why are we averse towards algorithms? A comprehensive literature review on algorithm aversion. In Proceedings of the 28th European Conference on Information Systems (ECIS), Marrakech, Morocco, 15–17 June 2020.
- Logg, J.M.; Minson, J.A.; Moore, D.A. Algorithm appreciation: People prefer algorithmic to human judgment. *Organ. Behav. Hum. Decis. Process.* **2019**, *151*, 90–103. [\[CrossRef\]](#)
- Bogert, E.; Schechter, A.; Watson, R.T. Humans rely more on algorithms than social influence as a task becomes more difficult. *Sci. Rep.* **2021**, *11*, 8028. [\[CrossRef\]](#) [\[PubMed\]](#)
- Longoni, C.; Cian, L. Artificial intelligence in utilitarian vs. hedonic contexts: The “word-of-machine” effect. *J. Mark.* **2022**, *86*, 91–108. [\[CrossRef\]](#)
- Cian, L.; Longoni, C.; Krishna, A. When disclosure of artificial intelligence boosts persuasion. *J. Consum. Res.* **2024**, advance online publication.
- Bettman, J.R.; Luce, M.F.; Payne, J.W. Constructive consumer choice processes. *J. Consum. Res.* **1998**, *25*, 187–217. [\[CrossRef\]](#)
- Holbrook, M.B.; Hirschman, E.C. The experiential aspects of consumption: Consumer fantasies, feelings, and fun. *J. Consum. Res.* **1982**, *9*, 132–140. [\[CrossRef\]](#)
- Chernev, A. When more is less and less is more: The role of ideal point availability and assortment in consumer choice. *J. Consum. Res.* **2003**, *30*, 170–183. [\[CrossRef\]](#)
- Chernev, A.; Böckenholt, U.; Goodman, J. Choice overload: A conceptual review and meta-analysis. *J. Consum. Psychol.* **2015**, *25*, 333–358. [\[CrossRef\]](#)
- Iyengar, S.S.; Lepper, M.R. When choice is demotivating: Can one desire too much of a good thing? *J. Personal. Soc. Psychol.* **2000**, *79*, 995–1006. [\[CrossRef\]](#) [\[PubMed\]](#)
- Xiao, B.; Benbasat, I. E-commerce product recommendation agents: Use, characteristics, and impact. *Mis Q.* **2007**, *31*, 137–209. [\[CrossRef\]](#)
- Fiske, S.T.; Cuddy, A.J.C.; Glick, P. Universal dimensions of social cognition: Warmth and competence. *Trends Cogn. Sci.* **2007**, *11*, 77–83. [\[CrossRef\]](#) [\[PubMed\]](#)

32. Epley, N.; Waytz, A. Mind perception. In *Handbook of Social Psychology*, 5th ed.; Fiske, S.T., Gilbert, D.T., Lindzey, G., Eds.; John Wiley & Sons: Hoboken, NJ, USA, 2010; pp. 498–541.
33. Hong, Y.; Cao, L.; Yang, M. Algorithmic labelling and consumer response in e-commerce. *Decis. Support Syst.* 2022, *advance online publication*.
34. Komiak, S.Y.X.; Benbasat, I. The effects of personalization and familiarity on trust and adoption of recommendation agents. *Mis Q.* 2006, 30, 941–960. [[CrossRef](#)]
35. Bergner, A.S.; Hildebrand, C.; Häubl, G. Machine talk: How verbal embodiment in conversational AI shapes consumer–brand relationships. *J. Consum. Res.* 2023, 50, 742–764. [[CrossRef](#)]
36. Kim, Y.; Sundar, S.S. Anthropomorphism of computers: Is it mindful or mindless? *Comput. Hum. Behav.* 2012, 28, 241–250. [[CrossRef](#)]
37. Sundar, S.S. The MAIN model: A heuristic approach to understanding technology effects on credibility. In *Digital Media, Youth, and Credibility*; Metzger, M.J., Flanagin, A.J., Eds.; The MIT Press: Cambridge, MA, USA, 2008; pp. 73–100.
38. Friestad, M.; Wright, P. The persuasion knowledge model: How people cope with persuasion attempts. *J. Consum. Res.* 1994, 21, 1–31. [[CrossRef](#)] [[PubMed](#)]
39. Li, B.; Lai, E.Y.; Wang, X. From tools to agents: Meta-analytic insights into human acceptance of AI. *J. Mark.* 2025, *advance online publication*.
40. Voss, K.E.; Spangenberg, E.R.; Grohmann, B. Measuring the hedonic and utilitarian dimensions of consumer attitude. *J. Mark. Res.* 2003, 40, 310–320. [[CrossRef](#)]
41. Dodds, W.B.; Monroe, K.B.; Grewal, D. Effects of price, brand, and store information on buyers' product evaluations. *J. Mark. Res.* 1991, 28, 307–319. [[CrossRef](#)]
42. MacKenzie, S.B.; Lutz, R.J. An empirical examination of the structural antecedents of attitude toward the ad in an advertising pretesting context. *J. Mark.* 1989, 53, 48–65. [[CrossRef](#)]
43. Tian, K.T.; Bearden, W.O.; Hunter, G.L. Consumers' need for uniqueness: Scale development and validation. *J. Consum. Res.* 2001, 28, 50–66. [[CrossRef](#)]
44. Polit, D.F.; Beck, C.T. The content validity index: Are you sure you know what's being reported? Critique and recommendations. *Res. Nurs. Health* 2006, 29, 489–497. [[CrossRef](#)] [[PubMed](#)]
45. Henseler, J.; Ringle, C.M.; Sarstedt, M. A new criterion for assessing discriminant validity in variance-based structural equation modeling. *J. Acad. Mark. Sci.* 2015, 43, 115–135.
46. Smithson, M. *Confidence Intervals; Quantitative Applications in the Social Sciences*; Sage: Thousand Oaks, CA, USA, 2003; Volume 140.
47. Fritz, M.S.; MacKinnon, D.P. Required sample size to detect the mediated effect. *Psychol. Sci.* 2007, 18, 233–239. [[CrossRef](#)] [[PubMed](#)]
48. Hayes, A.F. *Introduction to Mediation, Moderation, and Conditional Process Analysis*, 3rd ed.; The Guilford Press: New York, NY, USA, 2022.
49. MacKinnon, D.P.; Krull, J.L.; Lockwood, C.M. Equivalence of the mediation, confounding, and suppression effect. *Prev. Sci.* 2000, 1, 173–181. [[CrossRef](#)] [[PubMed](#)]
50. Faul, F.; Erdfelder, E.; Lang, A.-G.; Buchner, A. G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behav. Res. Methods* 2007, 39, 175–191. [[CrossRef](#)] [[PubMed](#)]
51. Hayes, A.F. An index and test of linear moderated mediation. *Multivar. Behav. Res.* 2015, 50, 1–22. [[CrossRef](#)]
52. Oppenheimer, D.M.; Meyvis, T.; Davidenko, N. Instructional manipulation checks: Detecting satisficing to increase statistical power. *J. Exp. Soc. Psychol.* 2009, 45, 867–872. [[CrossRef](#)]
53. Waytz, A.; Heafner, J.; Epley, N. The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *J. Exp. Soc. Psychol.* 2014, 52, 113–117. [[CrossRef](#)]
54. Brislin, R.W. Back-translation for cross-cultural research. *J. Cross-Cult. Psychol.* 1970, 1, 185–216. [[CrossRef](#)]
55. Beaton, D.E.; Bombardier, C.; Guillemin, F.; Ferraz, M.B. Guidelines for the process of cross-cultural adaptation of self-report measures. *Spine* 2000, 25, 3186–3191. [[CrossRef](#)] [[PubMed](#)]
56. DeVon, H.A.; Block, M.E.; Moyle-Wright, P.; Ernst, D.M.; Hayden, S.J.; Lazzara, D.J.; Savoy, S.M.; Kostas-Polston, E. A psychometric toolbox for testing validity and reliability. *J. Nurs. Scholarsh.* 2007, 39, 155–164. [[CrossRef](#)] [[PubMed](#)]
57. Lynn, M.R. Determination and quantification of content validity. *Nurs. Res.* 1986, 35, 382–385. [[CrossRef](#)]
58. Hu, L.; Bentler, P.M. Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Struct. Equ. Model.* 1999, 6, 1–55. [[CrossRef](#)]
59. Wang, C.L. New frontiers and future directions in interactive marketing: Inaugural editorial. *J. Res. Interact. Mark.* 2021, 15, 1–9. [[CrossRef](#)]
60. Wang, C.L. Editorial—What is an interactive marketing perspective and what are emerging research areas? *J. Res. Interact. Mark.* 2024, 18, 161–165. [[CrossRef](#)]

61. Boerman, S.C.; Willemsen, L.M.; Van Der Aa, E.P. “This post is sponsored”: Effects of sponsorship disclosure on persuasion knowledge and electronic word of mouth in the context of Facebook. *J. Interact. Mark.* **2017**, *38*, 82–92. [[CrossRef](#)]
62. Podsakoff, P.M.; MacKenzie, S.B.; Lee, J.-Y.; Podsakoff, N.P. Common method biases in behavioral research: A critical review of the literature and recommended remedies. *J. Appl. Psychol.* **2003**, *88*, 879–903. [[CrossRef](#)] [[PubMed](#)]
63. Bartneck, C.; Kulić, D.; Croft, E.; Zoghbi, S. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *Int. J. Soc. Robot.* **2009**, *1*, 71–81. [[CrossRef](#)]
64. Gray, K.; Knobe, J.; Sheskin, M.; Bloom, P.; Wegner, D.M. More than a body: Mind perception and the nature of objectification. *J. Personal. Soc. Psychol.* **2011**, *101*, 1207–1220. [[CrossRef](#)] [[PubMed](#)]
65. De Bruyn, A.; Viswanathan, V.; Beh, Y.S.; Brock, J.K.; Von Wangenheim, F. Artificial intelligence and marketing: Pitfalls and opportunities. *J. Interact. Mark.* **2020**, *51*, 91–105. [[CrossRef](#)]
66. De Freitas, J.; Nave, G.; Puntoni, S. Ideation with generative AI—In consumer research and beyond. *J. Consum. Res.* **2025**, *52*, 18–31. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.