



Article

Why Process-Based Explanations Foster Algorithmic Trust: A Procedural Justice Account of E-Commerce Recommendations

Ru Guo ¹, Bolu Wei ² and Xuemeng Guo ^{3,*}

¹ Experimental College, Open University of China, Beijing 100039, China; guor@ouchn.edu.cn

² China Academy of Safety Science and Technology, Beijing 100012, China; blwei1991@163.com

³ School of Economics and Management, Beijing Jiaotong University, Beijing 100044, China

* Correspondence: xmguo@bjtu.edu.cn

Abstract

E-commerce platforms increasingly rely on recommendation systems whose internal logic is often opaque, making explanation design important for consumer evaluation. Drawing on procedural justice theory, this study examines whether process-based explanations function as procedural justice cues in e-commerce recommendations and how they relate to algorithmic trust and continuance intention. In a between-subjects online experiment with 394 Chinese consumers (197 per condition), participants received either an outcome-based recommendation or a process-disclosure package that disclosed data inputs and reasoning and therefore bundled procedural content with greater specificity and informational richness. Relative to outcome-based explanations, this package increased perceived procedural justice and was associated with higher trust in the algorithm and greater continuance intention. Perceived procedural justice and trust formed a theoretically ordered indirect pathway, but this ordering should be read as theory-grounded rather than causally established because the mediators and outcome were measured contemporaneously. Exploratory moderation analyses suggested that responsiveness to process-based explanations reflected broader self-reported digital interpretive capacity rather than algorithm-specific literacy alone. Robustness checks further indicated that the procedural justice pathway was not eliminated by explanation clarity, cognitive load, scenario realism, product attractiveness, or privacy intrusiveness. The findings position process-disclosure packages as practical transparency tools while cautioning that their benefits depend on consumers' interpretive capacity and processing costs.

Keywords: algorithmic recommendation; process-based explanation; procedural justice; algorithmic trust; perceived algorithm literacy; continuance intention; human–AI interaction



Academic Editor: Qingfeng Zeng

Received: 5 May 2026

Revised: 21 June 2026

Accepted: 27 June 2026

Published: 1 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 \(CC BY\) license](#).

1. Introduction

E-commerce platforms increasingly rely on algorithmic recommendations to reduce search costs and personalize product discovery. These systems help consumers navigate dense online marketplaces [1,2] and increasingly mediate consumer relationships with firms [3]. Yet they also create a legitimacy problem: consumers respond not only to what is recommended but also to how it was produced—whether the platform used appropriate information and applied a recognizable procedure [4,5]. When recommendation procedures appear arbitrary or opaque, repeated reliance becomes harder to sustain.

Existing research on algorithmic transparency in digital commerce has examined explanations as mechanisms that reduce opacity [2,6], bundled explainable artificial intelligence (XAI) stimuli affecting satisfaction and trust [7], and artificial intelligence (AI)

literacy as an antecedent of consumer responses to AI-generated content [8,9]. These streams establish that transparency matters, but they often operationalize explainability as a broad system perception, a bundled feature set, or an individual capability. Less is known about how a specific explanation format functions as a procedural justice cue in e-commerce recommendations.

This distinction matters because explanation is not a single design choice. A platform may merely present a recommended item without explaining why it was selected. Alternatively, it may disclose the behavioral cues, product attributes, or decision rules used to generate the recommendation [2,5]. These formats invite different evaluative processes. An outcome-based explanation keeps attention on the recommended result, whereas a process-based explanation reveals part of the procedure that generated that result. The latter may allow consumers to judge whether the recommendation process appears consistent, relatively unbiased, and based on accurate information, all of which are central criteria in procedural justice theory [10–12]. Thus, the theoretical question is not simply whether transparency is beneficial, but whether explanation format provides consumers with the procedural information needed to evaluate algorithmic fairness.

To address this question, this study draws on procedural justice theory. Procedural justice research argues that individuals evaluate not only outcomes but also the procedures producing them [10,11]. Procedures are perceived as fair when they are applied consistently, suppress bias, use accurate information, allow correction, represent affected parties, and conform to ethical standards [10–12]. Although this framework was developed in human and organizational decision contexts, it is increasingly relevant to algorithmic decision systems, where people form fairness and trust judgments despite limited visibility into the system's inner workings [4,5]. Consumers typically observe interface cues rather than full algorithmic operations. Explanation format therefore becomes a practical mechanism through which consumers infer whether an algorithmic procedure warrants reliance.

The present study conceptualizes process-based explanation as a procedural justice cue in e-commerce recommendations. Rather than treating transparency as a broad system attribute, the study compares two explanation formats: outcome-based explanations, which present only the recommended product, and process-based explanations, which describe the information and reasoning used by the recommendation system. The theoretical claim is that process-based explanations may make the recommendation procedure easier for consumers to evaluate: they give consumers diagnostic material for judging whether the recommendation process appears consistent, relatively unbiased, and grounded in accurate information [2,5,10]. This study directly tests the effect of explanation type on perceived procedural justice; procedural evaluability is used as a conceptual interpretation of why that effect arises and is not itself measured. Perceived procedural justice is then expected to be positively associated with trust in the algorithm, and trust is expected to correspond to consumers' willingness to continue using the platform's recommendation feature [12–14].

Procedural evaluability refers to the diagnostic capacity of explanation content to make an otherwise opaque recommendation procedure assessable. It is used here as a conceptual bridge rather than as a separately measured mediator: it names the condition under which explanation content becomes relevant to procedural justice judgments, but it is not operationalized as an independent construct in this study. High transparency does not by itself imply high procedural evaluability; a platform may disclose extensive logs, model parameters, or technical detail that carries little diagnostic value for a consumer judging whether the recommendation procedure is fair. So defined, procedural evaluability extends and applies existing transparency and explainability ideas to the procedural-justice question rather than introducing a new measured construct.

The study also examines perceived algorithm literacy as a boundary condition. In this study, perceived algorithm literacy refers to consumers' self-reported understanding of how algorithmic decision systems use data, behavioral traces, and system rules to produce outputs [15–17]. This construct differs from product-domain expertise and general digital literacy. A consumer may know a great deal about tea products without understanding how recommendation algorithms work; another may be comfortable using mobile apps without understanding how recommendations are generated. Perceived algorithm literacy is theorized to strengthen the effect of process-based explanation on procedural justice because consumers who report greater algorithmic understanding should be more responsive to data and reasoning cues as evidence about the recommendation procedure [2,15]. The argument concerns interpretation of procedural information rather than unconditional reliance on algorithmic recommendations. Prior work on algorithm aversion and appreciation shows that expertise and experience may reduce reliance on algorithms in some contexts [9,18,19]. The present account therefore treats perceived algorithm literacy as responsiveness to procedural information, not as a general preference for algorithmic advice.

This study makes three focused contributions. First, it advances e-commerce explainability research by moving from transparency as a broad system perception to explanation format as a concrete design lever, clarifying when explanation content may make algorithmic recommendation procedures easier for consumers to evaluate [2,7,20]. Second, it extends procedural justice theory to e-commerce recommendation systems by interpreting process-based explanation as a candidate procedural-evaluability cue—a conceptual account we develop and motivate rather than an independently measured construct: explanation matters when it gives consumers criteria for judging whether the recommendation procedure appears consistent, relatively unbiased, and grounded in accurate information [10,11]. Third, it provides exploratory evidence that consumers' self-reported capacity to interpret digital and algorithmic cues conditions how they use explanation content, extending prior adoption-oriented accounts of AI literacy [8,15,18,19]. The same process information appears more consequential when consumers report greater interpretive capacity. A specificity analysis, however, indicates that this boundary condition is carried by general digital literacy rather than by algorithm-specific literacy alone; we therefore frame it as a broad self-reported digital interpretive capacity and treat the algorithm-specific moderation as not robust. An exploratory analysis further distinguishes these self-reported literacies from product-domain expertise, indicating that interpretive capacity and product-domain knowledge operate at different levels of consumer evaluation.

Figure 1 summarizes the conceptual model that guides the study. The model positions explanation type as the experimental treatment, perceived procedural justice and trust in the algorithm as sequential mediators, and perceived algorithm literacy as the first-stage boundary condition.

As shown in Figure 1, the model separates the mechanism from the boundary condition: process-based explanation is expected to increase perceived procedural justice, perceived procedural justice is theorized to be positively associated with algorithmic trust, and perceived algorithm literacy is theorized to moderate the strength of this first-stage path. This boundary condition is treated as exploratory, and Section 4.5.7 examines whether the moderation is specific to algorithm literacy or instead reflects a broader self-reported digital interpretive capacity.




Figure 1. Conceptual model of moderated serial mediation. The figure shows the hypothesized model; a specificity analysis (Section 4.5.7) indicates the first-stage moderation reflects general digital literacy rather than algorithm-specific literacy, so H5 and H6 are treated as exploratory.

2. Theoretical Background and Hypotheses

2.1. Algorithmic Explanations in E-Commerce Recommendations

Recommendation systems are central to contemporary e-commerce because they organize product information and reduce consumer search costs. These systems infer consumer preferences from browsing behavior, clicks, purchase histories, product attributes, and other platform data, then present selected options as recommendations [2,9,21,22]. Although such systems are often experienced as convenient, their underlying logic is largely hidden from consumers. The opacity of algorithmic recommendation raises a practical problem: consumers are asked to evaluate outputs without directly observing the procedure that produced them [5,15].

Explanations are one mechanism for reducing this opacity. Prior research on algorithmic transparency emphasizes that explanations can help users understand how algorithmic systems work, calibrate reliance, and evaluate whether system outputs are sensible, biased, or controllable [2,5,20]. In recommender systems, explanations may serve multiple purposes, including transparency, trust, effectiveness, persuasiveness, efficiency, satisfaction, and scrutability [1,2]. Yet explanations vary in content and depth. Some explanations merely identify the output; others reveal parts of the input data, reasoning logic, or rule structure used by the system [2,20].

This distinction is important for e-commerce recommendation design. An outcome-based explanation presents the recommended item without disclosing why it was selected. For example, a platform may display “Recommended for you” and show a product card. Such an explanation may be useful as an interface cue, but it gives consumers little basis for evaluating the underlying procedure. A process-based explanation goes further by describing the data inputs and reasoning logic behind the recommendation. For example, the platform may state that a product was recommended because the consumer recently viewed similar items, compared products across brands, and clicked products in a similar price range. This format does not fully disclose the algorithm, but it provides procedural cues that are absent from a simple outcome-based recommendation [2,5,23].

Much existing research treats transparency or explainability as a general perceived attribute. This approach is useful for identifying whether users value transparent systems, but it is less suited to explaining how particular explanation designs shape consumer evaluations [2,20]. Recent JTAER work has similarly examined XAI features as bundled

stimuli or service attributes that affect satisfaction, trust, purchase intention, or adoption trade-offs [7–9]. A process-based explanation, however, is not merely “more transparency.” It changes the object of consumer judgment. Instead of asking consumers to evaluate only the recommended product, it invites them to evaluate the process by which the product was selected. This shift provides the basis for applying procedural justice theory to algorithmic recommendation.

2.2. Procedural Justice and Algorithmic Recommendation

Procedural justice theory holds that people evaluate the fairness of decision-making procedures, not only the favorability of decision outcomes [10–12]. Leventhal’s formulation identifies several criteria that shape procedural fairness judgments, including consistency, bias suppression, accuracy, correctability, representativeness, and ethicality [10]. Colquitt’s construct validation later distinguished procedural justice from distributive, interpersonal, and informational justice, while also validating procedural rules such as consistency, freedom from bias, and accuracy as central components of fairness judgments [11]. Together, these works establish that procedures are evaluated as fair when they appear to apply rules consistently, avoid arbitrary or biased treatment, and rely on accurate information.

Algorithmic recommendations differ from traditional organizational decisions, but they raise a similar fairness problem. A recommendation is not a salary allocation or promotion decision, but it is still the result of a selection procedure. The platform chooses which product to display, which signals to weigh, and how to translate consumer behavior into a suggested item. Prior algorithmic decision research shows that people evaluate algorithmic systems through fairness, trust, emotion, objectivity, and legitimacy judgments, even when the systems are not human decision makers [4,5]. Consumers may not see the full algorithm, yet they can still form judgments about whether the process seems fair, reasonable, and grounded in appropriate information.

Drawing on Colquitt’s four-dimensional framework of organizational justice, this study treats procedural justice as especially suitable in this setting because the consumer’s central uncertainty concerns the recommendation process rather than the allocation of a scarce outcome. Distributive justice is less directly applicable because an e-commerce recommendation is not an allocation of wages, promotions, or benefits relative to inputs. Interpersonal justice is also less central because no human decision maker interacts with the consumer. Informational justice is relevant insofar as the platform provides an account of the recommendation, while the present focus is on whether that account enables evaluation of the underlying recommendation procedure [5,11]. Accordingly, the study focuses on perceived procedural justice, particularly the rule-application dimensions most applicable to single-shot algorithmic recommendations: consistency, bias suppression, and accuracy.

Explanation format is not treated here as procedural justice itself. Rather, process-based explanation is theorized as a diagnostic cue that enables consumers to evaluate whether the recommendation process satisfies core procedural criteria, especially consistency, bias suppression, and accuracy. When a platform explains that a recommendation is based on recent browsing, brand comparisons, similar price-range clicks, product ratings, and freshness information, consumers can infer whether the procedure uses relevant information and applies recognizable matching logic. These cues support judgments of accuracy and consistency. They may also reduce concerns that the recommendation is arbitrary or biased, because the platform provides a reasoned account of how the recommendation was generated [2,10].

Procedural evaluability therefore occupies a distinct position in the argument. Transparency describes the availability of information; explanation format describes how the platform organizes and presents recommendation reasons; procedural evaluability de-

scribes whether the available information is diagnostically useful for judging the legitimacy of the recommendation procedure; and perceived procedural justice captures the fairness judgment consumers form from those cues. Process-based explanations are expected to matter not because they disclose more information per se, but because they shift the consumer's evaluative target from the recommended object to the procedure that selected it.

Outcome-based explanations provide far less procedural information. They may still signal that the platform has selected something for the consumer, but they leave the process opaque. Consumers are therefore more likely to evaluate the recommendation as a bare output rather than as the result of an assessable procedure. Because procedural justice judgments require some basis for evaluating the decision process, outcome-based explanations should produce weaker procedural justice perceptions than process-based explanations [2,10,11].

H1. *Process-based explanations lead to higher perceived procedural justice than outcome-based explanations.*

2.3. Procedural Justice and Trust in the Algorithm

Trust is central to e-commerce because consumers often rely on platforms, vendors, and decision systems whose internal operations they cannot fully observe [13,21,24]. In algorithmic recommendation contexts, trust reflects confidence that the recommendation system is competent, reliable, and oriented toward users' interests, especially when product choice involves uncertainty, information overload, or difficulty assessing alternatives [20,21].

Procedural justice is theorized to be positively related to trust in the algorithm because fair procedures provide evidence that the system warrants reliance. In organizational justice research, fair procedures support trust by signaling that decision makers follow appropriate rules rather than acting arbitrarily or opportunistically [11,12]. A similar logic applies to algorithmic recommendation. When consumers perceive the recommendation process as consistent, relatively unbiased, and based on accurate information, they have stronger grounds for inferring that the algorithm is competent and reliable [4,5,20,25]. Explanations can shape trust and behavioral responses, but their effect depends on whether the explanation gives users a meaningful basis for evaluating the system [26].

H2. *Perceived procedural justice is positively related to trust in the recommendation algorithm.*

2.4. Trust in the Algorithm and Continuance Intention

Continuance intention refers to a consumer's willingness to keep using a system or feature after initial exposure. Information systems research distinguishes continuance from initial acceptance because platform value depends on sustained engagement rather than a one-time trial [14,24,27]. This distinction is especially relevant to e-commerce recommendation features, whose value depends on repeated interaction and ongoing reliance.

In algorithmic recommendation contexts, trust is the evaluation most proximal to continued use. Consumers who view the recommendation system as competent and reliable are more likely to use it again in future shopping, particularly when recommendations reduce search complexity and product uncertainty [13,21,24]. Related e-commerce research has linked algorithmic ethical perceptions and legitimacy to continuous usage intentions, underscoring the role of trust-related evaluations in sustained platform use [28].

In the present context, trust connects procedural evaluation to continued use. Perceived procedural justice provides evidence that the algorithmic process is legitimate, and greater trust in turn corresponds to stronger willingness to use the recommendation feature again [12,14].

H3. *Trust in the recommendation algorithm is positively related to continuance intention.*

The preceding arguments specify a theoretically ordered process: explanation format is expected to be linked to continuance intention by shaping how consumers evaluate both the recommendation procedure and the system that generates it, while leaving room for additional cue-based indirect paths. A serial mediation prediction follows from two considerations. First, continuance intention in e-commerce recommendation contexts is typically grounded in perceived legitimacy and reliability rather than in a one-time interface cue alone [13,14]. Second, procedural justice and trust are conceptually ordered: fairness perceptions concern whether procedures appear legitimate, whereas trust concerns whether the system warrants reliance. Procedural legitimacy provides evidence for trust formation [12,13], and trust is expected to correspond to willingness to continue using the recommendation feature. Process-based explanations should accordingly be linked to continuance through a serial pathway in which procedural justice and trust operate sequentially.

Taken together, H1–H3 specify an ordered process in which explanation format first shapes consumers' evaluation of the recommendation procedure. This procedural evaluation is expected to support trust in the algorithm, and trust is expected to correspond to continued willingness to use the recommendation feature. This logic implies not only separate path relationships but also a serial indirect pathway from explanation type to continuance intention through perceived procedural justice and algorithmic trust:

H4. *Perceived procedural justice and trust in the algorithm constitute a theoretically ordered serial indirect pathway linking explanation type to continuance intention.*

2.5. Self-Reported Digital Interpretive Capacity as an Exploratory Boundary Condition

Individuals differ in their capacity to make sense of digital and algorithmic interface cues. We refer to this broad, self-reported capacity as digital interpretive capacity, and we initially theorize one specific instantiation of it—perceived algorithm literacy—as the focal moderator. Because perceived algorithm literacy is conceptually proximate to general digital literacy, the boundary condition could in principle be carried either by an algorithm-specific capability or by a broader digital interpretive capacity. Section 4.5.7 tests this specificity directly, and the analysis ultimately favors the broader interpretation; we therefore retain perceived algorithm literacy as the originally hypothesized moderator below while flagging that the operative capacity appears to be broader.

Consumers vary in their self-reported understanding of algorithmic systems. Perceived algorithm literacy generally refers to users' knowledge of how algorithms use data, behavioral traces, and system rules to produce outputs [15]. In this study, the focal moderator is perceived algorithm literacy: consumers' reported understanding of how e-commerce recommendation systems use browsing, clicking, purchase patterns, product attributes, and system design choices to generate personalized suggestions. This perceived literacy is distinct from product-domain expertise and from general digital literacy. A consumer may be highly knowledgeable about tea products but know little about recommendation algorithms; another may be comfortable using mobile apps but still lack understanding of how platform recommendations are generated [15–17]. Conceptual distinctness does not, however, settle which capacity does the empirical work: perceived algorithm literacy may be separable in definition yet operate in practice as one facet of a broader digital interpretive capacity—a possibility we adjudicate directly in Section 4.5.7. Whereas recent JTAER research has examined AI literacy as a direct antecedent of consumer responses to AI-generated marketing content [8], this study theorizes perceived algorithm literacy as a boundary condition that shapes how consumers decode explanation format as a procedural justice cue.

The operative moderator is best understood as consumers' broader digital interpretive capacity, of which perceived algorithm literacy is the specific, algorithm-focused instantiation measured here; whether construed narrowly or broadly, this capacity should moderate the effect of explanation format on perceived procedural justice. Process-based explanations provide procedural information, but cues do not interpret themselves. Consumers need an interpretive schema that allows them to decode behavioral data, product attributes, and algorithmic reasoning as evidence of procedural rule compliance. Consumers reporting higher perceived algorithm literacy should be more responsive to the significance of data inputs and reasoning cues. When they read that a recommendation is based on prior browsing, brand comparisons, similar price-range clicks, ratings, and freshness information, they can more readily translate those cues into judgments about whether the system is using relevant and accurate information. They are therefore more likely to perceive process-based explanations as evidence of procedural fairness [2,15].

Leventhal's procedural justice framework offers a deeper reason for this moderation. Procedural justice judgments are not always equally activated; individuals are more likely to evaluate fairness when they have reason to suspect that a procedure could violate justice rules [10]. Consumers reporting higher perceived algorithm literacy may be more aware that algorithmic systems can be opaque, biased, or dependent on particular data traces [5,15,16]. This awareness should make them more attentive to procedural cues. For these consumers, a process-based explanation supplies the diagnostic material needed to assess whether the system has followed fair procedural rules. In contrast, consumers reporting lower perceived algorithm literacy may lack the interpretive schemas needed to translate behavioral and procedural cues into justice-relevant judgments. For them, a process-based explanation may appear as additional information to process rather than as evidence of consistency, bias suppression, or accuracy. This additional information can increase cognitive burden, consistent with cognitive load theory, and may make the recommendation process feel harder to evaluate [29], thereby reducing, but not necessarily reversing, the positive effect of process-based explanations on perceived procedural justice. Because the interpretive work this requires may draw on a general capacity to make sense of digital-interface cues, the moderating capability could be broader than algorithm-specific literacy; H5 and H6 accordingly retain perceived algorithm literacy as the originally specified, measurable indicator, while Section 4.5.7 tests directly whether the boundary is algorithm-specific or reflects this broader digital interpretive capacity.

This argument should be situated alongside prior research on algorithm aversion and algorithm appreciation. People may reject algorithmic advice after observing algorithmic errors, and algorithm appreciation may weaken among domain experts who prefer their own judgment [18,19,30]. Recent work on robo-advisors similarly suggests that transparency can be valued differently by high-expertise and general consumers [9]. These findings concern reliance, choice, or adoption behavior, often in relation to task or domain expertise. The present claim concerns interpretation of procedural cues: perceived algorithm literacy should strengthen the effect of process-based explanation on procedural justice because it reflects consumers' reported capacity to evaluate the recommendation procedure.

This distinction between perceived algorithm literacy and product-domain expertise is theoretically important. Domain experts may rely less on algorithms because they possess independent criteria for evaluating the recommended object [18,19]. Consumers reporting higher perceived algorithm literacy, by contrast, are expected to be more responsive to cues about whether the algorithmic procedure appears legitimate. Thus, perceived algorithm literacy, as the algorithm-focused indicator of this broader capacity, is expected to amplify the procedural-justice effect of process-based explanations while remaining conceptually distinct from product-domain reliance.

Given the self-reported nature of this moderator and its conceptual proximity to general digital literacy, H5 and H6 are tested as focal but exploratory boundary-condition hypotheses; Section 4.5.7 accordingly examines whether the first-stage moderation is specific to algorithm literacy or instead reflects a broader self-reported digital interpretive capacity.

Theoretically, the procedural cues disclosed in a process-based explanation—prior browsing, brand comparison, price-range clicks, ratings, and freshness information—are general digital-behavior concepts rather than technical properties of recommender algorithms. Decoding them as fairness-relevant evidence may therefore draw on broad familiarity with digital platforms rather than on specialized algorithmic knowledge. This reasoning anticipates the specificity result in Section 4.5.7: if the interpretive work the explanation demands is general rather than algorithm-specific, the moderation should be carried by general digital literacy, as we ultimately find.

H5. *Perceived algorithm literacy moderates the effect of explanation type on perceived procedural justice, such that the positive effect of process-based explanations is stronger among consumers reporting higher perceived algorithm literacy.*

If perceived algorithm literacy conditions the first-stage effect of explanation type on perceived procedural justice, it should also condition the broader indirect pathway. Process-based explanations should increase procedural justice more strongly among consumers reporting higher perceived literacy; stronger procedural justice perceptions would then be expected to correspond to greater trust in the algorithm, and this trust should be linked to higher continuance intention. Thus, the serial indirect effect of explanation type on continuance intention through procedural justice and trust should vary by perceived algorithm literacy. This proposition follows the logic of moderated mediation, in which the strength of an indirect effect depends on the level of a moderator [31].

This moderated serial mediation proposition integrates the mechanism and the boundary condition. It proposes that the downstream influence of process-based explanation depends on whether consumers can interpret procedural information as meaningful evidence of algorithmic fairness. For consumers reporting higher perceived algorithm literacy, process-based explanations should more strongly activate the procedural justice-trust-continuance pathway. For consumers reporting lower perceived algorithm literacy, the same explanation may generate a weaker indirect effect because procedural cues are less fully recognized, are processed as additional cognitive burden, or are not translated into fairness judgments [10,15,31].

H6. *Perceived algorithm literacy moderates the serial indirect effect of explanation type on continuance intention through perceived procedural justice and trust in the recommendation algorithm, such that the indirect effect is stronger among consumers reporting higher perceived algorithm literacy.*

3. Method

3.1. Research Design

This study used a two-cell between-subjects online experiment to examine how algorithmic explanation format affects consumers' perceived procedural justice, trust in the recommendation algorithm, and continuance intention. The independent variable was algorithmic explanation type, manipulated as either an outcome-based explanation or a process-based explanation. In the outcome-based condition, the platform presented the recommended product without explaining why it was selected. In the process-based condition, the platform explained the behavioral and product-related information used to generate the recommendation.

The experimental context was a mobile e-commerce recommendation scenario. Participants were asked to imagine that they were browsing for tea for daily consumption on a mobile e-commerce platform called ShopEase. The recommended product was held constant across conditions: “YunMist Daily Oolong Tea, 100 g.” The product category was selected because tea for daily consumption is familiar to Chinese consumers, relatively low in technical complexity, and unlikely to be naturally confounded with perceived algorithm literacy. The platform name was kept generic and non-branded to avoid contamination from existing platform attitudes.

Participants were assigned to one of two explanation conditions using the survey platform’s random equal-allocation function, which routes incoming respondents into the experimental conditions in approximately equal numbers. The final analytic sample contained 197 participants per condition. The outcome-based condition stated that ShopEase’s recommendation system selected the item for the participant and invited the participant to view price, reviews, delivery options, and detailed product information. The process-based condition included the same product and call-to-action and stated that the recommendation drew on recently viewed oolong teas, comparisons across different brands, clicks on tea products in a similar price range, product ratings, and freshness information. The process-based explanation deliberately excluded collaborative-filtering language such as “users similar to you” to avoid confounding explanation format with social-proof cues. In the English translations reported in Appendix A, the outcome-based and process-based messages contain approximately 30 and 57 words, respectively; these figures are reported to document the relative difference in explanation richness rather than to imply that the Chinese stimuli were length-matched. The full experimental stimuli are reported in Appendix A.

This manipulation operationalized process-based explanation as a procedurally diagnostic disclosure package. The design reflects a realistic platform contrast: outcome labels usually provide limited procedural information, whereas process explanations disclose additional cues about data inputs and reasoning logic. The process-based format is therefore accompanied by greater procedural specificity, textual detail, and informational richness. The analyses below evaluate whether the explanation-format effect remains evident after accounting for explanation clarity, cognitive load, perceived realism, product attractiveness, and privacy intrusiveness. The estimates should be interpreted as the effect of a practical process-disclosure design in which procedural orientation is accompanied by richer procedural content.

3.2. Participants and Data Collection

Data were collected through Wenjuanxing, a professional online survey platform in China. Respondents were recruited from the platform’s online participant pool and were eligible if they were adult consumers with prior e-commerce shopping experience. The platform implemented equal allocation across the two experimental conditions and applied standard quality-screening procedures during data collection. The research team further inspected the delivered responses for completeness, response time, attention-check performance, patterned responding, and response consistency. The sample was used to support experimental identification in an e-commerce consumer population and should not be interpreted as a nationally representative probability sample.

A total of 417 responses were received. Twenty-three responses were excluded based on the following predefined quality-screening criteria: duplicate submissions flagged by the survey platform from repeated device or IP signals ($n = 4$), failed attention checks ($n = 4$), reverse-item inconsistency ($n = 2$), excessively short completion time of less than 90 s ($n = 5$), responses below 80% completion ($n = 3$), and straight-lining ($n = 5$). These exclusions were

applied before any hypothesis testing and without reference to the substantive outcome variables. The final analytic sample comprised 394 valid responses, with 197 participants in the outcome-based explanation condition and 197 participants in the process-based explanation condition.

The final sample included 185 male respondents (47.0%) and 209 female respondents (53.0%). Respondents ranged in age from 20 to 55 years ($M = 36.34$, $SD = 9.81$). The age distribution was 121 respondents aged 20–29 years, 136 aged 30–39 years, 83 aged 40–49 years, and 54 aged 50–55 years. The two conditions did not differ significantly in age, Welch's $t = 1.32$, $p = 0.19$, or gender distribution, $\chi^2(1) = 1.47$, $p = 0.23$, consistent with the two conditions being demographically comparable. Full distributions for gender, age, education, and shopping frequency are reported in Table 1.

A two-sided sensitivity analysis for equal-sized group comparisons indicated that, with 197 participants per condition, $\alpha = 0.05$, and 80% power, the design could detect an effect of approximately Cohen's $d = 0.28$. The observed H1 effect ($d = 0.72$) was substantially larger than this threshold. Interaction and conditional indirect effects were evaluated using effect estimates and bootstrap confidence intervals.

Table 1. Demographic characteristics of the final analytic sample ($N = 394$).

Characteristic	Category	n	%
Gender	Male	185	47.0
	Female	209	53.0
Age (years)	20–29	121	30.7
	30–39	136	34.5
	40–49	83	21.1
	50–55	54	13.7
	Missing or unspecified	0	0.0
Education	High school or below	111	28.2
	Junior college	96	24.4
	Bachelor's degree	151	38.3
	Master's degree or above	33	8.4
	Missing or unspecified	3	0.8
Shopping frequency	Multiple times per week	114	28.9
	Once per week	107	27.2
	One to three times per month	97	24.6
	Less than once per month	74	18.8
	Missing or unspecified	2	0.5

Notes. $N = 394$. Percentages are rounded to one decimal place and may not sum to 100% due to rounding and missing or unspecified responses.

3.3. Measures

The questionnaire was administered in Chinese. Measurement items adapted from English-language sources were translated into Chinese and reviewed by bilingual researchers to ensure semantic equivalence in the e-commerce recommendation context. The English wording reported in this manuscript is a translation of the Chinese items used in the study. All constructs were measured using seven-point Likert scales ranging from 1 = strongly disagree to 7 = strongly agree. Unless otherwise noted, scale scores were computed by averaging their respective items. Full item wording is reported in Appendix B, and the survey flow and screening summary are reported in Appendix C.

Manipulation check. Four items assessed whether participants perceived the explanation as providing process information. Participants rated whether the recommendation explained why the product was selected, described the information used by the recommendation system, helped them understand how the recommendation was generated, and mainly showed a product outcome without explaining the selection process. The fourth item was reverse-coded for interpretation.

Perceived procedural justice. Perceived procedural justice was measured as context-specific perceived procedural justice in a single recommendation encounter, using six items adapted from Colquitt [11] and tailored to the algorithmic recommendation context. The items captured participants' perceptions that the recommendation process was fair, used relevant information, applied consistent matching principles, was relatively unbiased, was based on accurate information about shopping needs, and was justifiable. The measure primarily captured core procedural fairness criteria relevant to single-shot algorithmic recommendations, especially consistency, bias suppression, and accuracy. These dimensions were selected because they are most directly inferable from a single recommendation explanation; correctability and representativeness are better examined in settings that allow revision, voice, or repeated interaction. Consistent with Colquitt's measurement logic, this study used a Leventhal-criteria-based indirect measure that assessed specific procedural criteria rather than relying solely on a single global fairness judgment, which made the measure more diagnostic of the procedural elements underlying consumers' fairness evaluations [11].

Trust in the algorithm. Trust in the recommendation algorithm was initially measured with five author-adapted items conceptually grounded in IT trust research [13,21,24]. During item-level measurement diagnostics conducted before hypothesis testing, one item referring to confidence in relying on the recommendation system when shopping online was removed because it showed the lowest standardized loading and overlapped conceptually with downstream reliance and continuance intention. The retained four-item scale assessed participants' general trust in the recommendation system, belief in its ability to provide useful product suggestions, perceived reliability, and belief that the system acts in a way that benefits users like them. The originally administered item is reported transparently in Appendix B, and the before-and-after diagnostics for the trust item-deletion decision are reported in Appendix D, Table A8. The retained four items capture the intended dimensions of trust represented in the administered scale, including general trust, perceived capability and usefulness, perceived reliability, and the belief that the system acts in users' interests.

Continuance intention. Continuance intention was measured with three items adapted from Bhattacharjee [14] and modified for the single-shot e-commerce recommendation context. The items assessed participants' intention to continue using product recommendations on the platform, willingness to rely on platform recommendations in future shopping, and intention to keep using personalized recommendations.

Perceived algorithm literacy. Perceived algorithm literacy was measured with five author-developed items grounded in the algorithm awareness and literacy literature [15–17]. Because no established scale directly matched the e-commerce recommendation mechanism examined in the experiment, items were developed to capture consumers' perceived understanding of behavioral data use, recommendation generation, and system design choices in platform recommendations. The measure captures subjective or perceived algorithm literacy and is not designed as an objective technical-knowledge test. The items assessed participants' understanding that online recommendation systems use behavioral data, browsing, clicking, and purchase patterns, similar-user data, and system design choices to generate recommendations, as well as their ability to infer why platforms recommend particular products. The similar-user item captured general algorithm-awareness knowledge and was not part of the experimental stimulus, which deliberately omitted similar-user language to avoid confounding process explanation with social-proof cues. Perceived algorithm literacy was measured after the experimental task and is therefore treated as a measured individual-difference indicator in this experiment; as reported below, it did not differ significantly by condition.

Control and diagnostic variables. The study also measured perceived realism, cognitive load, privacy intrusiveness, product attractiveness, explanation clarity, tea domain expertise, and general digital literacy. Tea domain expertise and general digital literacy were included to distinguish perceived algorithm literacy from product-domain knowledge and general digital literacy. Privacy intrusiveness and explanation clarity were used in robustness analyses to evaluate the distinctiveness of the hypothesized procedural justice pathway relative to privacy concern and clearer wording.

3.4. Analytical Strategy

Explanation type was coded as 0 (outcome-based) and 1 (process-based). Manipulation-check items were first examined using Welch's independent-samples tests. Internal consistency was assessed using Cronbach's alpha for multi-item constructs. To strengthen measurement assessment, a confirmatory measurement model was estimated for the five focal multi-item constructs used in the main analyses: perceived procedural justice, trust in the algorithm, continuance intention, perceived algorithm literacy, and privacy intrusiveness. Standardized loadings, composite reliability (CR), average variance extracted (AVE), and heterotrait–monotrait (HTMT) ratios were inspected. To evaluate discriminant separation among perceived algorithm literacy, tea domain expertise, and general digital literacy, their pairwise correlations were also examined. All analyses were conducted in R version 4.5.3 using reproducible local analysis scripts; Welch tests and linear regressions were estimated in R, and indirect-effect intervals were obtained from nonparametric bootstrap resampling.

H1 was tested with an independent-samples comparison of perceived procedural justice across the two explanation conditions. H2 and H3 were assessed with linear regression models regressing trust in the algorithm on perceived procedural justice and continuance intention on trust in the algorithm, respectively. H4 was tested using a theoretically ordered serial mediation model in which explanation type was linked to continuance intention through perceived procedural justice and trust in the algorithm. For H5, perceived procedural justice was regressed on explanation type, mean-centered perceived algorithm literacy, and their interaction, with conditional effects probed at low, mean, and high levels of perceived algorithm literacy following recommendations for interpreting interaction effects [32]. H6 was tested using a moderated serial mediation model in which perceived algorithm literacy moderated the first-stage path from explanation type to perceived procedural justice. Indirect effects and the index of moderated serial mediation were estimated using 5000 nonparametric bootstrap resamples with 95% confidence intervals [33]. Unless otherwise stated, regression coefficients are reported as unstandardized coefficients (B). Robustness analyses examined whether the effect of explanation type on procedural justice remained evident after controls for explanation clarity, cognitive load, perceived realism, product attractiveness, and privacy intrusiveness. A supplemental reverse-order specification was also estimated, using the same bootstrapping procedure [33], to clarify that the measured mediator ordering is theory-grounded and statistically consistent with the focal model. These analyses serve as supplemental checks of whether the procedural justice pathway is empirically distinct from clearer wording, information load, scenario realism differences, product halo, privacy-related concern, and alternative ordering of the measured mediators.

Because several psychological constructs were measured in the same survey session, common method variance (CMV) remains a relevant consideration for associations among perceived procedural justice, trust, and continuance intention. The design mitigates this concern for the main treatment test because the independent variable was experimentally manipulated rather than self-reported, and the treatment effect on procedural justice

does not depend on a same-source predictor. The experimental design supports causal inference for the effect of explanation type on perceived procedural justice, whereas the ordering among procedural justice, trust, and continuance intention is theory-grounded and statistically consistent with the proposed model.

4. Results

4.1. Manipulation Check

The manipulation check confirmed that participants perceived the two explanation formats as intended. Compared with participants in the outcome-based condition, participants in the process-based condition reported significantly higher agreement that the recommendation explained why the product was selected ($M_{\text{outcome}} = 3.53$, $M_{\text{process}} = 4.39$, Welch's $t = 7.12$, $p < 0.001$, $d = 0.72$), described the information used by the platform's recommendation system ($M_{\text{outcome}} = 3.68$, $M_{\text{process}} = 4.60$, Welch's $t = 8.29$, $p < 0.001$, $d = 0.84$), and helped them understand how the platform generated the recommendation ($M_{\text{outcome}} = 3.52$, $M_{\text{process}} = 4.46$, Welch's $t = 8.17$, $p < 0.001$, $d = 0.82$). For the reverse-coded item, participants in the process-based condition reported significantly lower agreement that the recommendation mainly showed a product outcome without explaining the selection process ($M_{\text{outcome}} = 4.54$, $M_{\text{process}} = 3.59$, Welch's $t = -7.76$, $p < 0.001$, $d = -0.78$). All four manipulation-check items showed the expected direction and were statistically significant, indicating that the explanation-format manipulation was successful.

Figure 2 provides a visual summary of the manipulation check. For ease of interpretation, the fourth item is reverse-scored so that higher values consistently indicate stronger perceived process information. Across all four items, participants in the process-based condition perceived the recommendation as providing more process information than participants in the outcome-based condition.




Figure 2. Manipulation check: perceived process information. Bars represent group means; error bars indicate 95% confidence intervals.

The consistent separation across items supports the validity of the experimental manipulation before testing the hypothesized effects on procedural justice, trust, and continuance intention.

4.2. Reliability and Discriminant Checks

The focal theoretical constructs achieved acceptable internal consistency. Cronbach's alpha was 0.855 for perceived procedural justice, 0.797 for trust in the algorithm, 0.732 for continuance intention, 0.822 for perceived algorithm literacy, and 0.764 for privacy intrusiveness. The reliability of the three-item continuance intention scale was lower than that

of the other core constructs but remained above the conventional 0.70 threshold, consistent with the sensitivity of Cronbach's alpha to scale length and item intercorrelations [34]. Descriptive statistics and reliabilities are reported in Table 2; the correlation matrix is reported in Table 3. A five-factor confirmatory factor analysis (CFA) measurement model covering perceived procedural justice, trust in the algorithm, continuance intention, perceived algorithm literacy, and privacy intrusiveness showed good fit, $\chi^2(179) = 230.59$, $p = 0.006$, comparative fit index (CFI) = 0.981, Tucker–Lewis index (TLI) = 0.977, root mean square error of approximation (RMSEA) = 0.027, and standardized root mean square residual (SRMR) = 0.041. Table 4 reports loading ranges, composite reliability (CR), and average variance extracted (AVE); item-level standardized loadings are reported in Appendix D.

Table 2. Descriptive statistics and reliabilities.

Variable	M	SD	α
1. Perceived procedural justice	4.34	1.05	0.855
2. Trust in algorithm	4.14	1.03	0.797
3. Continuance intention	4.18	1.09	0.732
4. Perceived algorithm literacy	3.88	1.57	0.822
5. Perceived realism	4.42	0.83	0.600
6. Cognitive load	3.54	1.25	0.680
7. Privacy intrusiveness	3.67	1.18	0.764
8. Product attractiveness	4.41	0.95	0.650
9. Explanation clarity	4.01	1.05	0.690
10. Tea domain expertise	3.47	1.79	0.860
11. General digital literacy	3.79	1.41	0.820

Table 3. Correlation matrix.

Variable	1	2	3	4	5	6	7	8	9	10	11
1. Perceived procedural justice	—										
2. Trust in algorithm	0.36 ***	—									
3. Continuance intention	0.30 ***	0.33 ***	—								
4. Perceived algorithm literacy	−0.03	−0.01	0.07	—							
5. Perceived realism	0.11 *	−0.06	−0.01	−0.01	—						
6. Cognitive load	−0.08	−0.03	−0.10 *	−0.22 ***	−0.02	—					
7. Privacy intrusiveness	−0.24 ***	−0.27 ***	−0.25 ***	−0.07	0.01	0.28 ***	—				
8. Product attractiveness	0.28 ***	0.34 ***	0.13 *	0.05	0.01	−0.04	−0.12 *	—			
9. Explanation clarity	0.32 ***	0.29 ***	0.27 ***	−0.06	−0.02	0.09	−0.08	0.18 ***	—		
10. Tea domain expertise	−0.07	−0.05	−0.20 ***	−0.03	−0.03	−0.00	0.05	0.03	−0.05	—	
11. General digital literacy	−0.02	−0.02	0.01	0.34 ***	0.06	−0.20 ***	−0.02	0.05	0.07	−0.10 *	—

Notes. $N = 394$. Cronbach's alpha values are reported for multi-item constructs. Correlations are Pearson correlations. Several diagnostic variables were measured with short two-item scales and were used only as controls or robustness indicators. * $p < 0.05$, *** $p < 0.001$.

Several AVE values were slightly below the 0.50 benchmark, while CR values exceeded 0.70 and standardized loadings were substantively meaningful. Perceived algorithm literacy in particular was author-developed for this study and should be regarded as an exploratory measure pending external validation. Given the context-adapted nature of these measures and the brevity of several scales, the measures were judged adequate for hypothesis testing, while the constructs are interpreted with appropriate caution. Because below-threshold AVE indicates weaker-than-ideal convergent validity for some adapted scales, the findings should be read as evidence for the proposed empirical pattern rather than as definitive validation of every adapted measure. The pattern is nonetheless supported by CR values above 0.70, substantively meaningful loadings, HTMT ratios below conservative thresholds, and model results that were stable across alternative trust-scale specifications (Appendix D, Tables A8 and A9). The substantive conclusions were

unchanged in a robustness specification that retained the originally administered trust item later excluded from the focal four-item scale; Appendix D, Table A8 provides the before-and-after diagnostics for the trust item-deletion decision. Several short diagnostic scales—perceived realism ($\alpha = 0.60$), cognitive load ($\alpha = 0.68$), product attractiveness ($\alpha = 0.65$), and explanation clarity ($\alpha = 0.69$)—showed lower internal consistency, consistent with their two-item formats. Because these variables were used as diagnostics or robustness controls rather than focal theoretical constructs, they were retained as averaged indicators and interpreted with caution.

Table 4. Confirmatory measurement model diagnostics.

Construct	Standardized Loading Range	CR	AVE
Perceived procedural justice	0.664–0.739	0.855	0.497
Trust in algorithm	0.674–0.799	0.800	0.502
Continuance intention	0.637–0.730	0.731	0.476
Perceived algorithm literacy	0.666–0.708	0.822	0.480
Privacy intrusiveness	0.659–0.767	0.766	0.522

Notes. CR = composite reliability; AVE = average variance extracted. Item-level standardized loadings are reported in Appendix D. Trust was estimated using the retained four-item scale (TR1, TR2, TR3, TR5).

Discriminant separation among the focal constructs was acceptable. HTMT ratios were all below the conservative 0.85 benchmark, and the full HTMT matrix is reported in Appendix D, Table A7. Perceived algorithm literacy was also empirically distinguishable from both tea domain expertise and general digital literacy. The correlation between perceived algorithm literacy and tea domain expertise was negligible ($r = -0.034$), and the correlation between perceived algorithm literacy and general digital literacy was moderate ($r = 0.343$). The HTMT value between perceived algorithm literacy and general digital literacy was 0.417, below the 0.85 threshold, supporting their empirical separation despite their moderate zero-order correlation. The correlation between tea domain expertise and general digital literacy was also small ($r = -0.102$). A supplemental three-factor CFA for perceived algorithm literacy, tea domain expertise, and general digital literacy fit well, $\chi^2(21) = 22.75$, CFI = 0.999, TLI = 0.998, RMSEA = 0.015, and SRMR = 0.023, whereas a single-factor alternative fit poorly, $\chi^2(26) = 573.72$, CFI = 0.559, TLI = 0.390, RMSEA = 0.227, and SRMR = 0.161. These results support the conceptual distinction between perceived algorithm literacy, product-domain expertise, and general digital literacy.

4.3. Descriptive Statistics

Participants in the process-based explanation condition reported higher perceived procedural justice ($M = 4.70$, $SD = 0.96$) than those in the outcome-based condition ($M = 3.99$, $SD = 1.02$). They also reported higher trust in the algorithm ($M = 4.45$, $SD = 0.97$) than those in the outcome-based condition ($M = 3.83$, $SD = 1.01$), and higher continuance intention ($M = 4.36$, $SD = 1.05$) than those in the outcome-based condition ($M = 3.99$, $SD = 1.09$). Perceived algorithm literacy was comparable across conditions ($M_{\text{outcome}} = 3.94$, $SD = 1.53$; $M_{\text{process}} = 3.81$, $SD = 1.60$), suggesting that the moderator was not confounded with treatment assignment.

Process-based explanations imposed greater cognitive load than outcome-based explanations ($M_{\text{outcome}} = 3.17$, $M_{\text{process}} = 3.92$, Welch's $t = 6.19$, $p < 0.001$, $d = 0.62$), reflecting the greater informational content of process-based explanations. Privacy intrusiveness was nearly identical across conditions ($M_{\text{outcome}} = 3.67$, $M_{\text{process}} = 3.67$, Welch's $t = -0.04$, $p = 0.966$), indicating that the process-based explanation did not increase perceived privacy intrusiveness. Tea domain expertise ($M = 3.47$, $SD = 1.79$) and general digital literacy

($M = 3.79$, $SD = 1.41$) were at moderate levels and were empirically distinguishable from perceived algorithm literacy.

4.4. Hypothesis Testing

The hypothesis tests follow the conceptual model presented in Figure 1: H1–H3 examine the component direct paths, H4 examines the serial indirect pathway, H5 examines first-stage moderation, and H6 examines the resulting moderated serial mediation. All tests were conducted using the final analytic sample of 394 respondents. Results are summarized in Table 5, and the mediation and moderated mediation estimates are reported in Table 6.

Table 5. Summary of hypothesis tests.

Hypothesis	Path or Test	Estimate	Evidence	Conclusion
H1	Explanation type → Perceived procedural justice	$d = 0.72$	Welch's $t = 7.15$, $p < 0.001$	Supported
H2	Perceived procedural justice → Trust in algorithm	$B = 0.358$	$SE = 0.046$, $p < 0.001$, $R^2 = 0.132$	Supported
H3	Trust in algorithm → Continuance intention	$B = 0.344$	$SE = 0.050$, $p < 0.001$, $R^2 = 0.107$	Supported
H4	Theoretically ordered serial indirect pathway: Explanation type → PJ → TR → CI	indirect = 0.053	95% bootstrap CI [0.025, 0.090]	Supported
H5	Explanation type × Perceived algorithm literacy → Procedural justice	$B = 0.243$	$SE = 0.063$, $p < 0.001$, 95% CI [0.120, 0.366]	Supported in focal model; not algorithm-specific
H6	Index of moderated serial mediation	Index = 0.0182	95% bootstrap CI [0.0069, 0.0340]	Supported in focal model; not robust to general digital literacy

Table 6. Indirect effects and conditional process estimates.

Effect	Estimate	95% Bootstrap CI	Interpretation
Serial specific indirect effect: explanation type → PJ → TR → CI	0.053	[0.025, 0.090]	Significant
Specific indirect effect through PJ only	0.148	[0.065, 0.241]	Significant
Specific indirect effect through TR only	0.106	[0.045, 0.183]	Significant
Direct effect: explanation type → CI	0.063	[−0.142, 0.281]	Not significant
Index of moderated serial mediation (focal)	0.0182	[0.0069, 0.0340]	Significant (focal model)
Index after adding general digital literacy as a moderator	0.0078	[−0.0018, 0.0200]	Not significant

4.4.1. H1: Effect of Explanation Type on Perceived Procedural Justice

Hypothesis 1 predicted that process-based explanations would lead to higher perceived procedural justice than outcome-based explanations. This hypothesis was supported. Participants in the process-based explanation condition reported significantly higher perceived procedural justice ($M = 4.70$, $SD = 0.96$) than participants in the outcome-based explanation condition ($M = 3.99$, $SD = 1.02$), Welch's $t = 7.15$, $p < 0.001$, Cohen's $d = 0.72$. Thus, process-based explanations increased perceived procedural justice relative to outcome-based explanations. This interpretation applies within the bundled process-disclosure manipulation used here.

Figure 3 visualizes the main effect tested in H1. The figure plots group means and uncertainty intervals for perceived procedural justice, with individual responses overlaid to show the distribution within each experimental condition.




Figure 3. Main effect of explanation type on perceived procedural justice. Bars represent group means; error bars indicate 95% confidence intervals; individual responses are overlaid with horizontal jitter to show distribution.

The higher mean in the process-based condition is consistent with the procedural justice argument: disclosing the recommendation process gives consumers a stronger basis for evaluating the procedure as fair and reasonable.

4.4.2. H2: Procedural Justice and Trust in the Algorithm

Hypothesis 2 predicted that perceived procedural justice would be positively related to trust in the recommendation algorithm. This hypothesis was supported. Perceived procedural justice was positively related to trust in the algorithm, $B = 0.358$, $SE = 0.046$, $p < 0.001$, $R^2 = 0.132$. Consumers who perceived the recommendation process as more procedurally fair reported greater trust in the recommendation algorithm.

4.4.3. H3: Trust in the Algorithm and Continuance Intention

Hypothesis 3 predicted that trust in the recommendation algorithm would be positively related to continuance intention. This hypothesis was supported. Trust in the algorithm was positively related to continuance intention, $B = 0.344$, $SE = 0.050$, $p < 0.001$, $R^2 = 0.107$. Consumers who trusted the recommendation system more were more willing to continue using the platform's recommendation feature.

4.4.4. H4: Serial Mediation

Hypothesis 4 predicted a theoretically ordered serial indirect pathway from explanation type to continuance intention through perceived procedural justice and trust in the recommendation algorithm. The serial mediation analysis supported this pathway. The serial specific indirect effect of explanation type on continuance intention through perceived procedural justice and trust was significant, indirect effect = 0.053, 95% bootstrap CI [0.025, 0.090]. The direct effect of explanation type on continuance intention was not significant, direct effect = 0.063, 95% bootstrap CI [−0.142, 0.281].

The two additional specific indirect effects were also significant: the specific indirect effect through perceived procedural justice alone was 0.148, 95% bootstrap CI [0.065, 0.241], and the specific indirect effect through trust alone was 0.106, 95% bootstrap CI [0.045, 0.183].

These findings are consistent with the proposed psychological pathway and indicate that explanation format provides multiple convergent cues, with the hypothesized serial pathway operating alongside additional specific indirect effects via procedural justice alone and via trust alone.

4.4.5. H5: Moderating Role of Perceived Algorithm Literacy (Focal Model)

Hypothesis 5 predicted that perceived algorithm literacy would moderate the effect of explanation type on perceived procedural justice, such that the positive effect of process-based explanation would be stronger among consumers reporting higher perceived algorithm literacy. In a focal model that included only perceived algorithm literacy as a moderator, this hypothesis was supported: the interaction between explanation type and perceived algorithm literacy was positive and significant, $B = 0.243$, $SE = 0.063$, $p < 0.001$, 95% CI [0.120, 0.366]. As the specificity analysis in the Robustness section shows, however, this moderation is not specific to algorithm literacy and is carried by general digital literacy; H5 should therefore be read as exploratory.

Figure 4 plots the interaction between explanation type and perceived algorithm literacy. The gap between the process-based and outcome-based conditions widens as perceived algorithm literacy increases, indicating that consumers reporting higher perceived literacy were especially responsive to process information.




Figure 4. Moderating effect of perceived algorithm literacy. Lines represent model-predicted values; shaded bands indicate 95% confidence intervals.

In the focal model, this pattern is consistent with H5; the specificity analysis below shows, however, that the moderation reflects broad digital interpretive capacity rather than algorithm-specific literacy.

Simple-slope analysis showed that the effect of process-based explanation on perceived procedural justice was positive at low perceived algorithm literacy (-1 SD), $B = 0.329$; stronger at the mean level of perceived algorithm literacy, $B = 0.709$; and strongest at high perceived algorithm literacy ($+1$ SD), $B = 1.089$. These results indicate that process-based explanations improved perceived procedural justice across literacy levels, but the effect was substantially stronger among consumers reporting higher perceived algorithm literacy.

To complement the interaction plot, Figure 5 reports the conditional effects of process-based explanation at low, mean, and high levels of perceived algorithm literacy. This point-estimate view makes the increasing strength of the treatment effect more explicit.




Figure 5. Conditional effects of process-based explanation by perceived algorithm literacy. Points represent conditional effects at -1 SD, mean, and $+1$ SD of perceived algorithm literacy; error bars indicate 95% confidence intervals.

The conditional effects were positive across all three literacy levels and became progressively stronger from low to high perceived algorithm literacy. This pattern indicates amplification rather than reversal: consumers reporting lower perceived algorithm literacy benefited less, while consumers reporting higher perceived algorithm literacy showed a stronger procedural-justice response to process-based explanation.

4.4.6. H6: Moderated Serial Mediation

Hypothesis 6 predicted that perceived algorithm literacy would moderate the serial indirect effect of explanation type on continuance intention through perceived procedural justice and trust in the recommendation algorithm. In the focal model, this hypothesis was supported: the index of moderated serial mediation was positive and significant, index = 0.0182, 95% bootstrap CI [0.0069, 0.0340]. However, when general digital literacy was also permitted to moderate the first-stage path, this index was no longer significant, index = 0.0078, 95% bootstrap CI [−0.0018, 0.0200]; the specificity analysis below indicates that the conditional pattern is carried by general digital literacy rather than by algorithm-specific literacy.

In the focal model, this pattern is consistent with the proposed conditional process model. As the specificity analysis shows, the moderation is best interpreted as reflecting a broad self-reported digital interpretive capacity rather than an algorithm-specific effect.

4.5. Robustness Checks

4.5.1. Explanation Clarity

The first robustness check incorporated explanation clarity into the procedural-justice model. Explanation clarity was positively associated with perceived procedural justice, $B = 0.188$, $SE = 0.055$, $p < 0.001$. Including explanation clarity reduced the effect of explanation type on procedural justice from $B = 0.711$ to $B = 0.502$, corresponding to a 29.4% attenuation. The explanation-type coefficient nevertheless remained substantial and statis-

tically significant, $B = 0.502$, $SE = 0.116$, $p < 0.001$. This pattern indicates that explanation clarity partially overlaps with, but does not subsume, the procedural justice pathway.

Figure 6 summarizes the robustness check for explanation clarity. The coefficient for explanation type is attenuated after controlling for clarity, but it remains positive and statistically meaningful.




Figure 6. Robustness check: explanation clarity control. Bars represent unstandardized regression coefficients; error bars indicate 95% confidence intervals.

The robustness pattern indicates that clarity is part of the explanation-format effect while leaving a distinct procedural justice component. Process-based explanations appear to matter because they provide procedurally diagnostic information, not only because they are easier to understand.

4.5.2. Cognitive Load, Information Amount, and Full-Control Specification

Given that the process-based condition contained more procedural detail, an additional specification included cognitive load as a covariate. Controlling for cognitive load alone did not weaken the explanation-format effect; the estimate rose from $B = 0.711$ to $B = 0.833$, $SE = 0.102$, $p < 0.001$. When explanation clarity and cognitive load were entered together, explanation type remained a significant predictor of perceived procedural justice, $B = 0.637$, $SE = 0.119$, $p < 0.001$. In a fuller specification that simultaneously controlled for explanation clarity, cognitive load, privacy intrusiveness, perceived realism, product attractiveness, age, and gender, explanation type remained robust, $B = 0.622$, $SE = 0.116$, $p < 0.001$. These checks indicate that the procedural justice effect is not fully accounted for by wording clarity, processing burden, scenario realism, product halo, or privacy-related concern.

4.5.3. Scenario Realism

Scenario realism was examined because the process-based message contained more procedural detail. Perceived realism was slightly lower in the process-based condition ($M = 4.30$) than in the outcome-based condition ($M = 4.53$), Welch's $t = -2.70$, $p = 0.007$, $d = -0.27$. Both condition means were above the scale midpoint, indicating acceptable realism. When perceived realism was included as a covariate, the effect of explanation type on perceived procedural justice increased from $B = 0.711$ to $B = 0.756$, $SE = 0.099$, $p < 0.001$. This suppression pattern indicates that the small realism difference did not artificially inflate the main effect.

4.5.4. Product Attractiveness and Product-Halo Check

Product attractiveness was included to examine whether evaluations of the recommended product contributed to the procedural-justice effect. The process-based condition produced a small but significant product-attractiveness halo ($M_{\text{outcome}} = 4.24$, $M_{\text{process}} = 4.57$), Welch's $t = 3.51$, $p = 0.001$, $d = 0.35$. Re-estimating H1 with product attractiveness as a covariate showed that the explanation-type effect on perceived procedural justice remained robust, $B = 0.629$, $SE = 0.098$, $p < 0.001$, corresponding to 11.5% attenuation relative to the uncontrolled estimate. In the full serial mediation model, product attractiveness did not independently predict continuance intention when explanation type, perceived procedural justice, and trust were controlled, $B = -0.026$, $SE = 0.058$, $p = 0.651$. Thus, although process-based explanation slightly improved evaluation of the recommended product, the procedural justice pathway was empirically distinct from a product-halo account.

4.5.5. Privacy Intrusiveness

Privacy intrusiveness was examined as a potential parallel concern. Explanation type did not significantly predict privacy intrusiveness, $B = -0.005$, $SE = 0.119$, $p = 0.966$, indicating that the process-based explanation was not associated with higher perceived privacy intrusiveness. Privacy intrusiveness was negatively associated with continuance intention when explanation type, perceived procedural justice, and trust were controlled, $B = -0.138$, $SE = 0.045$, $p = 0.002$. These results suggest that privacy concern is relevant to continuance intention in algorithmic recommendation contexts while remaining empirically separate from the observed explanation-format effect.

4.5.6. Alternative Mediator Ordering

Because perceived procedural justice, trust, and continuance intention were measured in the same survey session, a supplemental robustness check examined the reverse serial ordering from explanation type to trust, then perceived procedural justice, and then continuance intention. This reverse-order serial indirect effect was also statistically significant but smaller, indirect effect = 0.037, 95% bootstrap CI [0.015, 0.065]. The focal sequence is retained because procedural justice theory treats perceived procedural legitimacy as evidence for trust formation; the supplemental result underscores that statistical mediation should be interpreted in conjunction with theoretical ordering.

4.5.7. Specificity of the Moderation: Perceived Algorithm Literacy Versus General Digital Literacy

Because perceived algorithm literacy was moderately correlated with general digital literacy ($r = 0.343$), an additional analysis examined whether the first-stage moderation was specific to algorithm literacy or instead reflected a broader self-reported digital interpretive capacity. The perceived algorithm literacy interaction remained significant when general digital literacy was added as a covariate, $B = 0.243$, $SE = 0.063$, $p < 0.001$. However, when general digital literacy was also allowed to interact with explanation type, the perceived algorithm literacy interaction was no longer significant, $B = 0.104$, $SE = 0.064$, $p = 0.104$, whereas the general digital literacy interaction was positive and significant, $B = 0.450$, $SE = 0.070$, $p < 0.001$. A direct test confirmed that the two interaction coefficients differed significantly, $t = 3.15$, $p = 0.002$, indicating that the divergence is not merely the difference between a significant and a non-significant coefficient. Consistent with this pattern, the moderated serial mediation index for perceived algorithm literacy was no longer significant once general digital literacy was permitted to moderate the explanation effect, index = 0.0078, 95% bootstrap CI [-0.0018, 0.0200].

Taken together, these results indicate that the first-stage moderation is carried by general digital literacy rather than by algorithm-specific literacy. The boundary condition is therefore best interpreted as a broad self-reported capacity to interpret digital and algorithmic cues, and H5 and H6 are accordingly treated as exploratory rather than confirmatory. This interpretation is compatible with the discriminant evidence reported in Section 4.2: perceived algorithm literacy and general digital literacy remain empirically distinct constructs ($HTMT = 0.417$), even though the first-stage moderation they produce overlaps substantially.

4.5.8. Common Method Variance

Because the mediators and outcome were measured in the same survey session, two complementary diagnostics evaluated whether CMV plausibly accounts for the focal associations. First, Harman's single-factor test was conducted on all items composing perceived procedural justice, the retained four-item trust scale, continuance intention, and perceived algorithm literacy (18 items). The first unrotated factor explained 25.56% of the variance, well below the 50% threshold conventionally interpreted as evidence against severe single-factor CMV. Second, a marker-variable analysis [35] was performed using general digital literacy as the marker. General digital literacy was selected because its mean absolute correlation with the focal mediators and outcome was small (mean $|r| = 0.015$; r with PJ = -0.020 , with TR = -0.015 , with CI = 0.009), and because it is theoretically distinct from the procedural and algorithm-specific perceptions captured by the focal constructs. General digital literacy is retained as the marker on this basis: marker validity depends on its near-zero zero-order association with the focal constructs (mean $|r| = 0.015$), a condition it satisfies and that the other available diagnostic variables satisfy least well. Although the specificity analysis above indicates that general digital literacy moderates the first-stage path from explanation type to perceived procedural justice, a first-stage interaction is a distinct statistical quantity from a zero-order association; a variable can exhibit a near-zero main-effect correlation with the focal constructs while still conditioning the strength of an experimental effect on one of them, so this moderation does not compromise the marker assumption. The smallest observed marker correlation was used as the CMV estimate ($r_M = 0.009$). After applying the partial correction $r_{\text{corrected}} = (r_{\text{observed}} - r_M) / (1 - r_M)$, the focal correlations remained materially unchanged: PJ-TR shifted from 0.363 to 0.357, PJ-CI from 0.300 to 0.293, and TR-CI from 0.327 to 0.320. All focal pairs that were significant before correction remained significant at $p < 0.001$ after correction. Together, the two diagnostics indicate that CMV does not plausibly account for the procedural justice pathway estimated above.

4.6. Summary of Findings

Overall, the results supported the hypothesized procedural justice pathway (H1–H4). The process-disclosure package represented by the process-based condition increased perceived procedural justice relative to the outcome-based condition. Perceived procedural justice was positively associated with trust in the recommendation algorithm, and trust was positively associated with continuance intention. The serial specific indirect effect from explanation type to continuance intention through procedural justice and trust was significant. The first-stage moderation (H5) and the moderated serial mediation (H6) were supported in the focal model but, as the specificity analysis showed, were carried by general digital literacy rather than algorithm-specific literacy; H5 and H6 are therefore treated as exploratory. Robustness analyses further supported the distinctiveness of the procedural justice pathway after accounting for explanation clarity, cognitive load, scenario realism differences, product-attractiveness halo, full-control specifications, and privacy

intrusiveness. Because the manipulation bundled procedural orientation with richer and more specific content, these effects are attributed to a process-disclosure package rather than to procedural information in isolation.

5. Discussion

This study examined how explanation format shapes consumer responses to algorithmic recommendations in e-commerce. Moving beyond a general “more transparency is better” view, the study compared outcome-based and process-based explanations and theorized the process-based condition as a procedural justice cue that may make the recommendation procedure easier for consumers to evaluate. The findings are consistent with the proposed model. The process-disclosure package represented by the process-based condition increased perceived procedural justice relative to the outcome-based condition. Perceived procedural justice was positively associated with trust in the recommendation algorithm, and trust was positively associated with continuance intention. The theoretically ordered serial indirect effect from explanation format to continuance intention through procedural justice and trust was significant. In addition, an exploratory boundary condition emerged: consumers’ self-reported capacity to interpret digital and algorithmic cues conditioned the first-stage effect of explanation format on procedural justice, although a specificity analysis showed this moderation to be carried by general digital literacy rather than by algorithm-specific literacy. Throughout the discussion that follows, the treatment is interpreted as a process-disclosure package: the process-based condition bundled procedural orientation with greater specificity, informational richness, and explanation length, so the results are read as evidence for the effectiveness of this package rather than as isolating the unique causal effect of procedural content alone.

These results support the central premise of the study: consumers do not simply react to algorithmic recommendations as outputs; they also evaluate the procedure by which those outputs are generated. A process-based explanation gives consumers more than an item-level recommendation. It provides diagnostic cues about data inputs and decision logic, allowing them to judge whether the recommendation process appears consistent, relatively unbiased, and grounded in appropriate information. This procedural evaluation is associated with, and provides a theoretical basis for, algorithmic trust and continued willingness to use the recommendation feature.

5.1. Theoretical Implications

5.1.1. Explanation Format as a Procedural Justice Cue

The first contribution of this study is to reframe algorithmic explanation format as a procedural justice cue. Prior research on explainable AI and algorithmic transparency has shown that explanations can increase users’ awareness of how algorithmic systems work, improve interpretability, and affect trust and acceptance [2,20]. Recent e-commerce research has similarly examined explainable AI features such as transparency, fairness, interpretability, and trustworthiness as drivers of satisfaction and purchase intention [7,9]. While this work has established the importance of explainability in digital commerce, it has tended to treat transparency as an aggregate perceptual feature rather than asking how specific explanation formats shape consumer evaluations.

The present results suggest that explanation format matters because it changes what consumers are able to evaluate. Outcome-based explanations keep attention on the recommended product. Process-based explanations shift attention to the procedure behind the recommendation and may thereby make the procedure easier for consumers to evaluate. Procedural evaluability is offered as a conceptual account of why explanations differ in theoretical consequence: explanation content matters when it supplies criteria by which

consumers can judge the legitimacy of an otherwise opaque recommendation process. Because the construct was not measured in this study, this account is an interpretation rather than an empirically established mechanism. In this account, procedural evaluability clarifies the condition under which explanation content becomes useful for procedural justice judgment. Procedural justice theory therefore suggests that fairness perceptions emerge when explanations provide diagnostic information that enables individuals to evaluate whether a procedure satisfies fairness-relevant criteria, such as consistency, bias suppression, and accuracy [10,11]. In the present context, process-based explanations provided consumers with cues about the recommendation system's use of browsing behavior, product comparisons, price-range clicks, ratings, and freshness information. These cues gave consumers material for procedural judgments about relevance, consistency, lack of bias, and accuracy. An alternative reading should also be acknowledged. Classic work on the mere presence of reasons suggests that explanations containing a "because" structure can elicit compliance even when the substantive content carries little informational weight [36]. The present design cannot fully separate the diagnostic value of disclosed procedural content from the more general signaling effect of providing any reasoned account, since the outcome-based condition contained no explicit reason-giving structure. A stricter test would compare process-based explanations with content-matched non-procedural explanations holding reason-giving constant.

This finding extends procedural justice theory into the domain of e-commerce recommendation systems by clarifying the object of consumer evaluation. Algorithmic decision research has shown that users evaluate algorithmic systems through fairness- and trust-related judgments, and that explanations can influence perceptions of justice in automated decisions [4,5]. Unlike organizational reward allocation, promotion, or dispute resolution, a product recommendation is not a formal allocation of scarce resources. Yet the consumer still confronts a procedural question: why was this product selected for me, and was the process reasonable? The results suggest that procedural justice theory can be productively applied to algorithmic recommendation contexts when explanation is understood as an interface-level cue that makes procedural justice evaluation possible.

5.1.2. Explanation Format as a Design-Specific Lever

The second contribution lies in shifting the discussion from transparency as a general attribute to explanation format as a design-specific lever. Much recent work on algorithmic transparency, explainable AI, and AI-enabled e-commerce examines whether users perceive a system as transparent, interpretable, fair, or trustworthy [2,7,20]. Other research shows that transparency may not benefit all consumers uniformly, especially when consumers face trade-offs among transparency, personalization, control, privacy, and cognitive effort [9]. This study complements those streams by isolating one concrete design choice: whether the system explains only the recommendation outcome or also the process behind it.

The results show that this design distinction has meaningful downstream consequences. Process-based explanations produced higher procedural justice perceptions, which were associated with stronger trust and continuance intention in the tested model. This finding is useful because broad guidance to increase transparency gives platform designers limited direction about which information to provide. A process-based explanation offers a practical form of transparency: it can preserve proprietary algorithmic detail while giving consumers enough procedural information to evaluate whether the system's reasoning appears legitimate.

The findings also clarify the conditions under which process-based explanation is useful. The robustness analyses showed that explanation clarity partially overlaps with the explanation-format effect, but does not subsume it. Explanation clarity attenuated the effect

by approximately 29.4%, yet the residual effect of explanation type remained substantial and statistically significant. The effect also remained robust after separately controlling for cognitive load, and in a fuller specification that simultaneously controlled for explanation clarity, cognitive load, privacy intrusiveness, perceived realism, product attractiveness, age, and gender. This pattern suggests that clear wording and processing demands matter while leaving a distinct procedural justice pathway. Process-based explanations appear to matter in part because they reveal information that consumers can use to evaluate the recommendation procedure.

The results also reveal a design trade-off. Process-based explanations were associated with higher cognitive load than outcome-based explanations ($M_{\text{process}} = 3.92$ vs. $M_{\text{outcome}} = 3.17$), indicating that richer procedural information imposes additional processing demands. The important point is that this cognitive cost coexisted with a procedural benefit. The moderation results suggest that consumers' self-reported digital interpretive capacity shapes whether they convert this additional information into justice-relevant evaluation. In this sense, process-based explanations create both an opportunity and a processing demand: they provide more procedural information, and consumers vary in their interpretive capacity to use that information effectively. Process-based explanations are therefore not costless transparency: they shift evaluative work onto the consumer in exchange for procedural information.

5.1.3. Digital Interpretive Capacity as a Boundary Condition

The third contribution concerns consumer heterogeneity. In the focal model, perceived algorithm literacy moderated the effect of explanation format on perceived procedural justice, with the effect strongest among consumers reporting higher perceived literacy. A specificity analysis, however, showed that this moderation was not specific to algorithm literacy: when general digital literacy was also allowed to moderate the explanation effect, the algorithm literacy interaction was no longer significant and was significantly weaker than the general digital literacy interaction ($t = 3.15$, $p = 0.002$). We therefore interpret the boundary condition as a broad, self-reported capacity to interpret digital and algorithmic cues rather than as an algorithm-specific effect, and we treat H5 and H6 as exploratory. This pattern is consistent with the view that procedural cues require interpretive schemas: consumers must translate behavioral and system cues into fairness-relevant judgments, and consumers reporting greater digital interpretive capacity appear better equipped to do so.

This finding refines how perceived algorithm literacy should be understood in e-commerce research. Because the specificity analysis indicates that the operative moderator is a broad self-reported digital interpretive capacity rather than algorithm-specific literacy, the remarks below refer to perceived algorithm literacy as the measured indicator while treating the underlying boundary condition as digital interpretive capacity. Algorithmic knowledge is not evenly distributed among users, and users vary in their ability to understand how algorithmic systems select, prioritize, and personalize information [15]. Recent JTAER research has examined AI literacy as a direct antecedent of consumer responses to AI-generated marketing content [8]. The present study instead positions perceived algorithm literacy as a boundary condition that shapes how consumers decode explanation format. Consumers reporting higher perceived algorithm literacy appear more responsive to the procedural significance of data inputs and reasoning cues. They may more readily infer that a process-based explanation indicates the system is using relevant information and applying a reasoned matching logic. Consumers reporting lower perceived algorithm literacy, by contrast, may benefit less because the same explanation can appear as additional information to process rather than as evidence of consistency, bias suppression, or accuracy.

This result also clarifies the relationship between perceived algorithm literacy and prior work on algorithm aversion or appreciation. Prior research has shown that people may reject algorithmic advice after observing algorithmic errors, and that algorithm appreciation can weaken among domain experts who prefer their own judgment [18,19]. Recent work on robo-advisors also shows that transparency may be valued differently across consumer segments with different expertise levels [9]. The present finding is more specific: self-reported digital interpretive capacity increases responsiveness to procedural information. Reported understanding of digital and algorithmic systems does not necessarily mean unconditional acceptance of algorithmic recommendations; it helps consumers evaluate whether a particular algorithmic procedure appears legitimate.

This distinction also helps explain why process-based explanations did not produce a uniform effect. At low levels of perceived algorithm literacy, the effect of process-based explanation on procedural justice remained positive but smaller. At high levels of perceived algorithm literacy, the effect was much stronger. This pattern suggests that process-based explanations may be most effective when consumers have the cognitive schema needed to interpret behavioral traces, product attributes, and algorithmic reasoning as evidence of fair procedure. Explanation design and user literacy therefore operate jointly: platforms can disclose procedural information, but consumers must also be able to recognize its procedural significance. Why, then, did algorithm-specific literacy not retain unique explanatory power once general digital literacy was allowed to moderate the effect? The most plausible reading is that the explanation did not require algorithm-specific knowledge to be useful: its cues were everyday digital-behavior concepts interpretable through general platform experience, so the narrower construct added little beyond the broader one. This carries a direct implication for AI- and algorithm-literacy research. Self-reported algorithm literacy may behave less like a distinct technical competence than like one expression of a broader digital interpretive capacity, particularly in consumer-facing settings where explanations are deliberately simplified. Future work should separate algorithm-specific technical knowledge from general digital interpretive competence—ideally pairing self-reports with objective knowledge tests—before attributing consumer responses to algorithm literacy specifically.

5.1.4. Process-Based Explanations as Multi-Cue Trust Signals

The serial mediation results support the proposed procedural justice pathway, while the additional specific indirect effects suggest that process-based explanations may also carry broader trust signals. This pattern complements the procedural justice account: some cues help consumers evaluate whether the recommendation procedure appears fair, whereas others may more directly suggest system competence, reliability, or usefulness [13,21,24–26]. Procedural justice remains the focal theoretical route identified in this study; broader trust cues operate alongside that route rather than replacing it. Notably, the serial specific indirect effect through procedural justice and trust (0.053) was smaller in magnitude than the specific indirect effects operating through procedural justice alone (0.148) or trust alone (0.106). The serial pathway is therefore best read as one theoretically ordered route among several convergent cue-based pathways activated by process-based explanation, alongside parallel routes that operate independently of the strict serial sequence. This pattern is consistent with the procedural justice account but indicates that process-based explanation also conveys cues that bear on trust and continuance more directly than the strict serial sequence implies. Because procedural justice, trust, and continuance intention were measured within a single session, and because a reverse serial ordering (trust before procedural justice) was also statistically significant (Section 4.5.6), the procedural-justice-then-trust ordering is retained on theoretical rather than empirical grounds. The serial pathway should be read as theory-grounded and statistically consistent

with the data, not as a causally established sequence; the present design licenses a causal claim only for the experimentally manipulated effect of explanation format on perceived procedural justice.

5.1.5. Domain Expertise as an Object-Level Boundary

An exploratory boundary finding concerns tea domain expertise. Tea expertise was included primarily as a discriminant covariate, but it showed a small negative association with continuance intention, $r = -0.20$, $p < 0.001$. This association remained significant in a model controlling for explanation type, perceived procedural justice, trust in the algorithm, and general digital literacy, $B = -0.109$, $SE = 0.028$, $p < 0.001$. As tea expertise was not a focal experimental moderator, the finding is interpreted as exploratory.

The finding is consistent with algorithm-aversion and algorithm-appreciation research showing that expertise can reduce reliance on algorithmic advice when users prefer their own judgment or believe they can evaluate the domain independently [18,19]. It also clarifies the difference between perceived algorithm literacy and product-domain expertise. Perceived algorithm literacy concerns understanding of the recommendation procedure; product-domain expertise provides independent criteria for judging the recommended object. Thus, algorithm-level understanding may strengthen the use of process cues, whereas object-level expertise may reduce dependence on the system for product choice.

Future research could examine this two-layer structure more directly by manipulating perceived algorithm literacy and product-domain expertise in the same design.

5.2. Practical Implications

The findings translate into three design principles for personalized recommendation systems in familiar, low-to-moderate involvement product contexts. Each principle follows from the same design logic: process information is useful when it helps consumers evaluate the procedure behind a recommendation.

First, because procedural justice judgments depend on whether consumers can identify the basis of a recommendation, platforms can identify the procedural basis for the recommendation. At the default recommendation-card level, a platform can move beyond a bare “recommended for you” label by briefly identifying the data categories or product signals that informed the recommendation, such as recent browsing, comparable brands, similar price-range clicks, ratings, or freshness information. Such wording does not require revealing proprietary models, but it gives consumers a procedural basis for evaluating why the product appeared.

Second, because procedural fairness judgments depend on whether the information used appears relevant and accurate, platforms can explain the relevance of data signals. A useful process explanation clarifies why the named signals matter for the recommendation—for example, that recent browsing indicates category interest, comparable brands indicate preference boundaries, and similar price-range clicks indicate budget fit. This turns data disclosure into procedural justification rather than a mere inventory of tracked behavior.

Third, because the effect of process information depends on consumers’ self-reported digital interpretive capacity, platforms can layer explanation depth. A short card-level cue can provide minimal procedural evaluability without crowding the interface; a plain-language “Why am I seeing this?” panel can then explain data categories and matching logic; and expandable details can clarify how behavioral signals and product attributes are combined. This structure preserves transparency while avoiding the burden of forcing every user to process a dense explanation.

Concretely, and consistent with the broader transparency and explanation-design literature, which motivates matching disclosure depth to consumers' needs and constraints [2,9,20], layered disclosure can be operationalized as three escalating tiers tied to how much interpretive effort a consumer is willing to invest; Table 7 specifies each tier, its placement, example interface microcopy, and the users it best serves.

Table 7. A layered process-disclosure design for e-commerce recommendation explanations.

Disclosure Tier	Placement and Trigger	What It Discloses (Example Microcopy)	Best Suited for
Tier 1: card-level cue	Shown by default on the recommendation card	Names the procedural basis in one plain-language line, without technical terms: "Recommended because you recently viewed oolong teas and compared similar price-range products."	All users; keeps cognitive load minimal
Tier 2: "Why am I seeing this?" panel	Revealed on demand, one tap from the card	Justifies each named signal in everyday language, turning a data inventory into a procedural rationale: "We used your recent category browsing, product comparisons, and rating preferences to match this item."	Users who want a fuller rationale
Tier 3: expandable detail and controls	Revealed on demand, expanded from Tier 2	Lists the signal categories, indicates how they are combined, and offers controls to edit or remove signals: "Signals considered: recent browsing, brand comparison, price range, ratings, freshness. You can edit or remove these signals."	Users seeking mechanism-level detail; higher digital interpretive capacity

Notes. Platforms should default to the shallowest tier and let depth be user-initiated rather than exposing every user to a dense explanation. Because responsiveness to process information appears to vary with consumers' self-reported digital interpretive capacity, design choices can be tuned to the audience: for consumers reporting lower interpretive capacity, favor short, non-technical phrasing, concrete examples, and icons, and avoid surveillance-like wording; for those reporting higher capacity, Tiers 2–3 can expose signal types, approximate weighting, data provenance, and control options. In privacy-sensitive categories, surface the edit-or-remove control earlier to preserve autonomy. Because platforms cannot directly observe a user's interpretive capacity, explanation depth is best left user-initiated, with tier defaults informed by observable interaction behavior rather than by inferred literacy.

As a practical heuristic, then, the interface can disclose the procedural basis to everyone in one line, justify it on demand, and reserve mechanism-level detail for the users who opt in.

At the same time, process-based explanations are not universally beneficial, and several trade-offs temper the recommendations above. First, richer procedural disclosure increases cognitive load, which may burden rather than help consumers with lower digital interpretive capacity. Second, disclosing behavioral inputs can heighten privacy or surveillance concerns in more sensitive product categories, partly offsetting the procedural-justice benefit. Third, transparent disclosure of ranking signals may allow sophisticated users or sellers to game recommendation inputs strategically. Fourth, detailed explanations can reduce rather than increase trust when they reveal weak, overly commercial, or irrelevant recommendation logic. These trade-offs reinforce the case for layered disclosure: concise procedural cues by default, plain-language explanations on demand, and deeper detail only when users seek it. These principles are best implemented with sensitivity to consumer literacy, privacy, and autonomy. The findings apply most directly to moderate-sensitivity product contexts of the kind studied here; in higher-sensitivity categories such as financial, health, or insurance products, process-based explanation may need to balance procedural disclosure against privacy reactance. Consumers reporting higher self-reported

digital interpretive capacity may use richer process information to evaluate procedural legitimacy, whereas those reporting lower capacity may need simpler language, concrete examples, and less technical phrasing. Platform designers can avoid surveillance-like wording and make clear that explanation is meant to support consumer judgment, not pressure consumers to delegate choice to the algorithm. Product-domain expertise and perceived algorithm literacy should also be separated in interface design: consumers who know a product category well may still want to understand the recommendation procedure, while users reporting higher perceived algorithm literacy may use process explanations to evaluate legitimacy rather than to surrender product judgment [37].

5.3. Limitations and Future Research

The findings suggest several boundary conditions and directions for future research.

First, the study used a simulated e-commerce recommendation scenario rather than observing actual platform behavior. This design allowed for causal identification of the explanation-format effect on perceived procedural justice because participants were randomly assigned to outcome-based or process-based conditions. Future research could test whether the same procedural justice pathway predicts actual platform metrics such as recommendation-card click-through rates, dwell time on the recommendation module, add-to-cart decisions, conversion rates, feature retention, repeat recommendation use, or repeat purchase behavior in field experiments. The effects observed here may also vary when consumers interact with real platforms under time pressure, competing product options, and accumulated platform experience.

Second, the study focused on a single product category: oolong tea for daily consumption. This category was chosen because it is familiar to Chinese consumers and avoids strong technical-product confounds, so the findings generalize most directly to low-to-moderate involvement e-commerce recommendations involving familiar consumer goods. The procedural justice pathway documented here may unfold differently in at least four boundary categories that the present design did not test. (i) In high-involvement or high-risk decision categories, such as financial products, insurance, or health-related goods, correctability and user control may rise in salience and partly displace the consistency, bias-suppression, and accuracy cues emphasized here; process-based explanation may therefore need to disclose recourse options rather than only data inputs and reasoning logic. (ii) In privacy-sensitive categories, such as health, finance, or location-based services, disclosing behavioral inputs may simultaneously make the recommendation procedure easier for consumers to evaluate and heighten privacy concern; the procedural justice benefit observed in the present moderate-sensitivity context may be partially offset by privacy reactance. (iii) In credence-good categories, such as supplements, professional services, or expert advice, where consumers cannot independently verify product quality even after purchase, procedural cues may carry greater weight than in search-good categories like daily tea, potentially amplifying the procedural justice effect. (iv) In categories with substantial product-domain expertise dispersion, such as wine, photography equipment, or specialty foods, the object-level boundary identified in Section 5.1.5 may become more pronounced, with high-expertise consumers further discounting algorithmic recommendations regardless of explanation format. Future research could test whether the procedural justice pathway holds, weakens, or interacts with category-specific concerns across these boundary conditions, and whether category type itself moderates the optimal disclosure depth.

Third, future studies can separate process orientation from explanation richness more directly. This study examined a process-based explanation that disclosed individual behavioral cues and product-level information, deliberately excluding collaborative-filtering

language such as “users similar to you” to avoid social-proof contamination. The manipulation reflects a realistic disclosure package in which process orientation is accompanied by greater procedural detail, specificity, and textual content. A four-cell design crossing explanation orientation and explanation richness—brief outcome explanation, detailed outcome explanation, brief process explanation, and detailed process explanation—could determine whether procedural evaluability arises from process orientation itself, from the amount of procedural detail, or from their combination. Relatedly, the procedural justice scale includes items on the use of relevant information and information accuracy, which partially echo the data inputs named in the process-based stimulus; the H1 effect may therefore reflect measurement–stimulus content overlap in addition to a procedural fairness judgment. That said, this overlap is confined to two of the six procedural-justice items (relevant and accurate information); the remaining items—fairness, consistency, freedom from bias, and justifiability—do not restate the stimulus content yet shifted in the same direction, a pattern that a pure wording-overlap account would not predict. The overlap is therefore best read as a target for cleaner stimulus–scale separation in future work rather than as an alternative explanation of the present H1 effect. Relatedly, procedural evaluability was treated throughout as a conceptual bridge rather than a measured variable. A direct test would develop items capturing the perceived diagnosticity of explanation content for fairness judgments and examine whether procedural evaluability statistically mediates the effect of explanation type on perceived procedural justice; until then, the proposed mechanism remains a theoretical proposition awaiting operationalization.

Fourth, trust in the algorithm was measured as a four-item composite reflecting the trust content typically described in the algorithmic trust literature, including general trust, perceived capability, perceived reliability, and a benevolence-related belief that the system acts in users’ interests. This approach is consistent with much of the algorithmic trust literature, but it does not distinguish cognitive trust from emotional trust. Classical IS research suggests that cognitive trust and emotional trust may play different roles in recommendation-agent adoption, especially when users delegate decisions to a system rather than use it as a decision aid [21]. Future studies could separately measure cognitive and emotional trust to determine whether procedural justice primarily strengthens rational beliefs about system competence or emotional confidence in relying on the system.

Fifth, perceived algorithm literacy was measured as subjective understanding of recommendation mechanisms and was administered after the experimental task. The measure is conceptually grounded in algorithm awareness and literacy research but should be interpreted as perceived algorithm literacy rather than objective technical knowledge [15]. The moderation result therefore concerns perceived algorithmic understanding. The self-report measure also cannot separate perceived literacy from related constructs such as confidence in one’s own knowledge or technological self-efficacy. Perceived algorithm literacy did not differ significantly across experimental conditions, and future studies can measure literacy before exposure, combine subjective literacy with objective knowledge tests, include self-efficacy measures, or manipulate algorithmic literacy through educational interventions. Such designs would clarify how literacy shapes consumers’ interpretation of process-based explanations. More broadly, several focal measures should be read as relatively exploratory. Perceived algorithm literacy was author-developed for this context and has no prior psychometric track record, and several constructs returned average variance extracted below the 0.50 convergent-validity benchmark (perceived procedural justice 0.497, continuance intention 0.476, perceived algorithm literacy 0.480). The reliability, HTMT, and trust-scale stability checks reported in Section 4.2 mitigate but do not eliminate these construct-validity concerns; the corresponding estimates should therefore be interpreted

with caution, and replication with validated or objective measures is a precondition for firm conclusions.

Finally, future longitudinal or experimental mediation designs could further examine the temporal ordering among perceived procedural justice, trust, and continuance intention. The present experiment supports causal inference for the effect of explanation format on perceived procedural justice, while the downstream pathway is based on measured psychological constructs collected in the same survey session. Because the mediators and outcome were measured contemporaneously, alternative causal orderings among procedural justice, trust, and continuance intention cannot be ruled out; the proposed sequence should be read as theory-grounded rather than as an established causal chain. Future studies could manipulate procedural justice cues separately from explanation format or measure trust and continuance after repeated platform interactions.

6. Conclusions

This study investigated how algorithmic explanation format shapes consumer responses to e-commerce recommendations. The findings show that a process-disclosure package—process-based explanation combined with greater specificity and richer content—increased perceived procedural justice relative to outcome-based explanations. Perceived procedural justice and trust formed a theoretically ordered indirect pathway linking explanation format to continuance intention, with the ordering grounded in theory rather than established causally, given contemporaneous measurement of the mediators and outcome. Exploratory analyses further suggest that consumers' self-reported capacity to interpret digital and algorithmic cues conditions responsiveness to process-based explanations; this boundary condition, however, was carried by general digital literacy rather than algorithm-specific literacy and should be interpreted as a broad digital interpretive capacity.

The central implication is that algorithmic explanation is not merely a transparency feature. It is a procedural design cue. When platforms explain how a recommendation is generated, they give consumers material for evaluating whether the system's process appears consistent, relatively unbiased, and based on relevant information. These evaluations matter because they are closely tied to trust and continued engagement with recommendation systems.

For e-commerce platforms, the results suggest that explanation design can move beyond simple outcome labels. Process-disclosure packages can support procedural legitimacy without requiring full algorithmic disclosure. At the same time, explanation design should remain sensitive to user literacy and cognitive burden. A useful explanation is not necessarily the longest or most technical one, but the one that helps consumers understand enough of the process to judge whether the recommendation is fair and trustworthy.

Overall, this study advances e-commerce research by linking explanation format, perceived procedural justice, algorithmic trust, and continuance intention in a single conditional process model. It shows that the value of explainable recommendation systems depends on both what platforms disclose and whether that disclosure helps consumers evaluate the recommendation procedure as evidence of a fair process.

Author Contributions: Conceptualization, R.G. and B.W.; methodology, R.G.; software, R.G.; validation, R.G., B.W. and X.G.; formal analysis, R.G.; investigation, R.G.; resources, B.W. and X.G.; data curation, R.G.; writing—original draft preparation, R.G.; writing—review and editing, R.G., B.W. and X.G.; visualization, R.G.; supervision, B.W. and X.G.; project administration, R.G.; funding acquisition, R.G. and X.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Natural Science Foundation of China, grant number 72272010, and the Scientific Research Project of Open University of China, grant number Q23C0019.

Institutional Review Board Statement: Ethical review and approval were waived for this study due to its anonymous, minimal-risk online survey design involving adult participants, no intervention beyond exposure to low-risk e-commerce recommendation scenarios, and no collection, storage, or access by the research team to participants' IP addresses, device identifiers, or other personally identifiable information.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the corresponding author upon reasonable request, subject to anonymization and applicable institutional restrictions.

Acknowledgments: The authors thank Pan Liu for constructive suggestions and assistance during the revision of this manuscript.

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

Appendix A. Experimental Stimuli

The questionnaire was administered in Chinese; the English stimulus wording reported below is provided for transparency. The platform name, product, and call-to-action were held constant across conditions; only the explanation format varied.

Table A1. Common scenario shown to all participants.

Stimulus Text
Imagine that you are shopping on a mobile e-commerce platform called ShopEase. You are browsing for tea for daily consumption. After looking through several products, you see a product recommendation generated by the platform's recommendation system.

Table A2. Experimental explanation conditions.

Condition	Stimulus Text
Outcome-based explanation	Recommended for you: YunMist Daily Oolong Tea, 100 g. ShopEase's recommendation system selected this item for you. Tap the product card to view price, reviews, delivery options, and detailed product information.
Process-based explanation	Recommended for you: YunMist Daily Oolong Tea, 100 g. ShopEase's recommendation system selected this item because you recently viewed oolong teas, compared options across different brands, and clicked tea products in a similar price range. The system also considers product ratings and freshness information. Tap the product card to view price, reviews, delivery options, and detailed product information.

Design note. The process-based condition deliberately excluded collaborative-filtering wording such as "users similar to you" to avoid confounding explanation format with social-proof cues.

Appendix B. Measurement Items

All scale items used a seven-point Likert response format (1 = strongly disagree, 7 = strongly agree). The English wording below is a translation of the questionnaire administered to participants.

Table A3. Measurement items and questionnaire wording.

Construct/Variable	Code	Item Wording
Manipulation check	MC1	The recommendation explained why this product was selected for me.
Manipulation check	MC2	The recommendation described the information used by the platform's recommendation system.
Manipulation check	MC3	The recommendation helped me understand how the platform generated the product recommendation.
Manipulation check	MC4_R	The recommendation mainly showed me a product outcome without explaining the selection process. [Reverse-coded]
Perceived procedural justice	PJ1	The recommendation process seemed fair.
Perceived procedural justice	PJ2	The recommendation process appeared to use relevant information.
Perceived procedural justice	PJ3	The recommendation process seemed consistent in how it matched products with user preferences.
Perceived procedural justice	PJ4	The recommendation process seemed relatively unbiased.
Perceived procedural justice	PJ5	The recommendation process appeared to be based on accurate information about my shopping needs.
Perceived procedural justice	PJ6	The recommendation process seemed justifiable.
Trust in algorithm	TR1	I trust the platform's recommendation system.
Trust in algorithm	TR2	I believe the recommendation system is capable of providing useful product suggestions.
Trust in algorithm	TR3	I believe the recommendation system would make reliable recommendations.
Trust in algorithm	TR4	I feel confident relying on this recommendation system when shopping online. [Administered but excluded from the final trust scale after item-level diagnostics because of conceptual overlap with reliance and continuance intention.]
Trust in algorithm	TR5	The recommendation system seems to act in a way that benefits users like me.
Continuance intention	CI1	I would continue using product recommendations on this platform.
Continuance intention	CI2	I would be willing to rely on this platform's recommendations in future shopping.
Continuance intention	CI3	I intend to keep using personalized recommendations when shopping on this platform.
Perceived realism	REAL1	The shopping scenario felt realistic.
Perceived realism	REAL2	The recommendation screen was similar to what I might see on an actual e-commerce platform.
Cognitive load	COG1	Reading the recommendation required a lot of mental effort.
Cognitive load	COG2	The recommendation message was difficult to process.
Privacy intrusiveness	PI1	The recommendation made me concerned about how my data were used.
Privacy intrusiveness	PI2	The recommendation felt somewhat intrusive.
Privacy intrusiveness	PI3	The platform seemed to know more about my behavior than I expected.
Product attractiveness	PA1	The recommended product seemed attractive to me.
Product attractiveness	PA2	I would be interested in learning more about this product.
Explanation clarity	CL1	The recommendation message was easy to understand.
Explanation clarity	CL2	The explanation was clearly written.
Perceived algorithm literacy	AL1	I understand that online recommendation systems use user behavior data to generate suggestions.

Table A3. *Cont.*

Construct/Variable	Code	Item Wording
Perceived algorithm literacy	AL2	I know that recommendation algorithms may use browsing, clicking, and purchase patterns.
Perceived algorithm literacy	AL3	I understand that recommendations can be influenced by data from users with similar behavior.
Perceived algorithm literacy	AL4	I can usually infer why an e-commerce platform recommends certain products to me.
Perceived algorithm literacy	AL5	I understand that algorithmic recommendations are shaped by both user data and system design choices.
Tea domain expertise	TE1	I am knowledgeable about tea products.
Tea domain expertise	TE2	I can evaluate the quality of tea better than most consumers.
General digital literacy	DL1	I am comfortable using digital platforms and mobile apps.
General digital literacy	DL2	I can quickly learn how to use new online shopping features.
Attention check	IMC1	To confirm that you are reading carefully, please select “somewhat disagree” (i.e., 3) for this item.

Demographic variables included age, gender, highest education level, occupation type, and monthly e-commerce shopping frequency.

Appendix C. Survey Flow and Screening Summary

The online questionnaire was implemented in Wenjuanxing. Incoming respondents were randomly assigned to one of the two explanation conditions, and the condition code was recorded as 0 = outcome-based explanation and 1 = process-based explanation.

Table A4. Survey flow.

Step	Block	Content
1	Opening and informed consent	Welcome message, study purpose, anonymity statement, voluntary participation, and consent confirmation.
2	Scenario and random assignment	Common e-commerce scenario followed by random assignment to the outcome-based or process-based explanation vignette; condition code recorded.
3	Manipulation check	MC1-MC4_R measured perceived process information immediately after exposure to the vignette.
4	Focal constructs	Perceived procedural justice, trust in algorithm, and continuance intention.
5	Diagnostic and control variables	Perceived realism, cognitive load, privacy intrusiveness, product attractiveness, and explanation clarity.
6	Attention check	Instructional manipulation check asking participants to select the specified response option.
7	Moderator and discriminant covariates	Perceived algorithm literacy, tea domain expertise, and general digital literacy.
8	Demographics	Age, gender, education, occupation, and e-commerce shopping frequency.

Table A5. Quality-screening exclusions.

Screening Category	n	Note
Initial responses received	417	Raw responses delivered by the online survey platform.
Platform-flagged duplicate device or IP submissions	4	Flagged by the survey platform; excluded before hypothesis testing.
Failed attention checks	4	Excluded before hypothesis testing.
Reverse-item inconsistency	2	Excluded before hypothesis testing.
Completion time less than 90 seconds	5	Excluded before hypothesis testing.
Response below 80% completion	3	Excluded before hypothesis testing.
Straight-lining	5	Excluded before hypothesis testing.
Total excluded	23	All exclusions were applied without reference to substantive outcome variables.
Final analytic sample	394	Outcome-based explanation: $n = 197$; process-based explanation: $n = 197$.

These quality-screening rules were applied before hypothesis testing and without reference to the substantive outcome variables.

Appendix D. Measurement Model Diagnostics

Table A6. Item-level standardized loadings for focal measurement model.

Construct	Item	Standardized Loading
Perceived procedural justice	PJ1	0.739
Perceived procedural justice	PJ2	0.716
Perceived procedural justice	PJ3	0.691
Perceived procedural justice	PJ4	0.664
Perceived procedural justice	PJ5	0.721
Perceived procedural justice	PJ6	0.695
Trust in algorithm	TR1	0.674
Trust in algorithm	TR2	0.675
Trust in algorithm	TR3	0.677
Trust in algorithm	TR5	0.799
Continuance intention	CI1	0.700
Continuance intention	CI2	0.730
Continuance intention	CI3	0.637
Perceived algorithm literacy	AL1	0.703
Perceived algorithm literacy	AL2	0.681
Perceived algorithm literacy	AL3	0.707
Perceived algorithm literacy	AL4	0.708
Perceived algorithm literacy	AL5	0.666
Privacy intrusiveness	PI1	0.767
Privacy intrusiveness	PI2	0.738
Privacy intrusiveness	PI3	0.659

Notes. Standardized loadings are from the five-factor confirmatory measurement model reported in Section 4.2. TR4 was administered in the questionnaire but excluded from the focal trust scale and therefore is not included in the measurement model. Item wording is reported in Appendix B.

Table A7. Heterotrait–monotrait (HTMT) ratio matrix for focal and discriminant constructs.

Construct	PJ	TR	CI	AL	PI	DL	TE
PJ	—						
TR	0.439	—					
CI	0.381	0.429	—				
AL	0.088	0.061	0.096	—			
PI	0.298	0.350	0.331	0.116	—		
DL	0.058	0.055	0.039	0.417	0.046	—	
TE	0.082	0.089	0.260	0.044	0.064	0.122	—

Notes. HTMT = heterotrait–monotrait ratio of correlations. PJ = perceived procedural justice; TR = trust in algorithm; CI = continuance intention; AL = perceived algorithm literacy; PI = privacy intrusiveness; DL = general digital literacy; TE = tea domain expertise. Values are reported to three decimal places. All HTMT ratios were below the conservative 0.85 threshold. Item-deletion transparency check. TR4 was retained in a supplemental five-item trust composite as a robustness check. The substantive direction and statistical support for the focal trust-related conclusions were unchanged, but the four-item scale is retained because TR4 had the lowest loading and overlapped conceptually with downstream reliance and continuance intention.

Table A8. Trust scale before-and-after diagnostics for the TR4 item-deletion decision.

Scale Version	Cronbach's α	CR	AVE	Standardized Loading Range	TR4 Loading	Decision Note
Original five-item trust scale (TR1–TR5)	0.811	0.815	0.470	0.585–0.790	0.585	TR4 showed the lowest standardized loading and overlapped conceptually with downstream reliance and continuance intention.
Retained four-item trust scale (TR1, TR2, TR3, TR5)	0.797	0.800	0.502	0.674–0.799	—	Retained as the focal trust scale because it improved AVE while preserving the intended trust content.

Notes. The original five-item trust composite included TR1 through TR5. TR4 referred to confidence in relying on the recommendation system when shopping online. It was excluded from the focal trust scale because it showed the lowest standardized loading and overlapped conceptually with downstream reliance and continuance intention. The retained four-item scale comprises TR1, TR2, TR3, and TR5.

Table A9. Robustness check: serial mediation re-estimated with the original five-item trust scale.

Effect/Path	Estimate	95% Bootstrap CI/SE	Interpretation
PJ → five-item trust	$B = 0.364$	$SE = 0.045$	$p < 0.001$
Five-item trust → CI	$B = 0.357$	$SE = 0.052$	$p < 0.001$
Serial: explanation type → PJ → trust → CI	0.057	[0.027, 0.097]	Significant
Specific indirect through PJ only	0.144	[0.061, 0.239]	Significant
Specific indirect through trust only	0.097	[0.039, 0.173]	Significant
Direct effect: explanation type → CI	0.072	[−0.143, 0.288]	Not significant

Notes. The model re-estimates the focal serial mediation using the original five-item trust scale (TR1–TR5) with 5000 nonparametric bootstrap resamples. Substantive conclusions are unchanged relative to the focal four-item specification reported in Table 6. The five-item scale is reported only as a robustness check; the four-item scale is retained as focal because it yields higher convergent validity (AVE = 0.502 vs. 0.470; see Table A8).

References

1. Tintarev, N.; Masthoff, J. Explaining recommendations: Design and evaluation. In *Recommender Systems Handbook*, 2nd ed.; Ricci, F., Rokach, L., Shapira, B., Eds.; Springer: New York, NY, USA, 2015; pp. 353–382. [\[CrossRef\]](#)
2. Rader, E.; Cotter, K.; Cho, J. Explanations as mechanisms for supporting algorithmic transparency. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 21–26 April 2018*; Association for Computing Machinery: New York, NY, USA, 2018; Paper 103, pp. 1–13. [\[CrossRef\]](#)
3. Puntoni, S.; Reczek, R.W.; Giesler, M.; Botti, S. Consumers and artificial intelligence: An experiential perspective. *J. Mark.* **2021**, *85*, 131–151. [\[CrossRef\]](#)
4. Lee, M.K. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data Soc.* **2018**, *5*, 2053951718756684. [\[CrossRef\]](#)

5. Binns, R.; Van Kleek, M.; Veale, M.; Lyngs, U.; Zhao, J.; Shadbolt, N. 'It's reducing a human being to a percentage': Perceptions of justice in algorithmic decisions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 21–26 April 2018*; Association for Computing Machinery: New York, NY, USA, 2018; Paper 377, pp. 1–14. [\[CrossRef\]](#)
6. Haque, A.K.M.B.; Islam, A.K.M.N.; Mikalef, P. Explainable artificial intelligence (XAI) from a user perspective: A synthesis of prior literature and problematizing avenues for future research. *Technol. Forecast. Soc. Change* **2023**, *186*, 122120. [\[CrossRef\]](#)
7. Almamy, A.S.M.; Habib, S.; Nasser, L.K.; Hamadneh, N.N. Effect of explainable AI features on user satisfaction and purchase intention in Saudi mobile shopping apps. *J. Theor. Appl. Electron. Commer. Res.* **2026**, *21*, 120. [\[CrossRef\]](#)
8. Liu, Q.; Osman, L.H.; Lian, Z.; Wel, C.A.C.; Hamid, S.N.A. From perception to purchase: How AI literacy shapes consumer decisions in AI-generated sponsored vlogs across products and services. *J. Theor. Appl. Electron. Commer. Res.* **2025**, *20*, 302. [\[CrossRef\]](#)
9. Choi, J.; Kang, S.; Moon, J.; Jeon, S.; Lim, S. Algorithmic transparency and consumer trade-offs in AI-based financial e-commerce services. *J. Theor. Appl. Electron. Commer. Res.* **2026**, *21*, 86. [\[CrossRef\]](#)
10. Leventhal, G.S. What should be done with equity theory? New approaches to the study of fairness in social relationships. In *Social Exchange: Advances in Theory and Research*; Gergen, K.J., Greenberg, M.S., Willis, R.H., Eds.; Plenum Press: New York, NY, USA, 1980; pp. 27–55.
11. Colquitt, J.A. On the dimensionality of organizational justice: A construct validation of a measure. *J. Appl. Psychol.* **2001**, *86*, 386–400. [\[CrossRef\]](#) [\[PubMed\]](#)
12. Colquitt, J.A.; Conlon, D.E.; Wesson, M.J.; Porter, C.O.L.H.; Ng, K.Y. Justice at the millennium: A meta-analytic review of 25 years of organizational justice research. *J. Appl. Psychol.* **2001**, *86*, 425–445. [\[CrossRef\]](#) [\[PubMed\]](#)
13. McKnight, D.H.; Choudhury, V.; Kacmar, C. Developing and validating trust measures for e-commerce: An integrative typology. *Inf. Syst. Res.* **2002**, *13*, 334–359. [\[CrossRef\]](#)
14. Bhattacharjee, A. Understanding information systems continuance: An expectation-confirmation model. *MIS Q.* **2001**, *25*, 351–370. [\[CrossRef\]](#)
15. Cotter, K.; Reisdorf, B.C. Algorithmic knowledge gaps: A new dimension of (digital) inequality. *Int. J. Commun.* **2020**, *14*, 745–765.
16. Eslami, M.; Rickman, A.; Vaccaro, K.; Aleyasen, A.; Vuong, A.; Karahalios, K.; Hamilton, K.; Sandvig, C. 'I always assumed that I wasn't really that close to [her]': Reasoning about invisible algorithms in news feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, Seoul, Republic of Korea, 18–23 April 2015*; Association for Computing Machinery: New York, NY, USA, 2015; pp. 153–162. [\[CrossRef\]](#)
17. DeVito, M.A.; Birnholtz, J.; Hancock, J.T.; French, M.; Liu, S. How people form folk theories of social media feeds and what it means for how we study self-presentation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 21–26 April 2018*; Association for Computing Machinery: New York, NY, USA, 2018; Paper 120, pp. 1–12. [\[CrossRef\]](#)
18. 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. [\[CrossRef\]](#) [\[PubMed\]](#)
19. Logg, J.M.; Minson, J.A.; Moore, D.A. Algorithm appreciation: People prefer algorithmic to human judgment. *Organ. Behav. Hum. Decis. Processes* **2019**, *151*, 90–103. [\[CrossRef\]](#)
20. Shin, D. The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI. *Int. J. Hum.-Comput. Stud.* **2021**, *146*, 102551. [\[CrossRef\]](#)
21. 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. [\[CrossRef\]](#)
22. Zhang, Y.; Chen, X. Explainable recommendation: A survey and new perspectives. *Found. Trends Inf. Retr.* **2020**, *14*, 1–101. [\[CrossRef\]](#)
23. Chen, C.; Tian, A.D.; Jiang, R. When post hoc explanation knocks: Consumer responses to explainable AI recommendations. *J. Interact. Mark.* **2024**, *59*, 234–250. [\[CrossRef\]](#)
24. Gefen, D.; Karahanna, E.; Straub, D.W. Trust and TAM in online shopping: An integrated model. *MIS Q.* **2003**, *27*, 51–90. [\[CrossRef\]](#)
25. Ning, X.; Lu, Y.; Li, W.; Gupta, S. How transparency affects algorithmic advice utilization: The mediating roles of trusting beliefs. *Decis. Support Syst.* **2024**, *183*, 114273. [\[CrossRef\]](#)
26. Leichtmann, B.; Humer, C.; Hinterreiter, A.; Streit, M.; Mara, M. Effects of explainable artificial intelligence on trust and human behavior in a high-risk decision task. *Comput. Hum. Behav.* **2023**, *139*, 107539. [\[CrossRef\]](#)
27. Maslowska, E.; Malthouse, E.C.; Hollebeek, L.D. The role of recommender systems in fostering consumers' long-term platform engagement. *J. Serv. Manag.* **2022**, *33*, 721–732. [\[CrossRef\]](#)
28. Liu, Y.; Sun, X. Towards more legitimate algorithms: A model of algorithmic ethical perception, legitimacy, and continuous usage intentions of e-commerce platforms. *Comput. Hum. Behav.* **2024**, *150*, 108006. [\[CrossRef\]](#)
29. Sweller, J. Cognitive load during problem solving: Effects on learning. *Cogn. Sci.* **1988**, *12*, 257–285. [\[CrossRef\]](#) [\[PubMed\]](#)

30. Mahmud, H.; Islam, A.K.M.N.; Luo, X.R.; Mikalef, P. Decoding algorithm appreciation: Unveiling the impact of familiarity with algorithms, tasks, and algorithm performance. *Decis. Support Syst.* **2024**, *179*, 114168. [[CrossRef](#)]
31. Hayes, A.F. An index and test of linear moderated mediation. *Multivar. Behav. Res.* **2015**, *50*, 962683. [[CrossRef](#)] [[PubMed](#)]
32. Aguinis, H.; Edwards, J.R.; Bradley, K.J. Improving our understanding of moderation and mediation in strategic management research. *Organ. Res. Methods* **2017**, *20*, 665–685. [[CrossRef](#)]
33. Hayes, A.F. *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach*, 2nd ed.; Guilford Press: New York, NY, USA, 2018.
34. Cortina, J.M. What is coefficient alpha? An examination of theory and applications. *J. Appl. Psychol.* **1993**, *78*, 98–104. [[CrossRef](#)]
35. Lindell, M.K.; Whitney, D.J. Accounting for common method variance in cross-sectional research designs. *J. Appl. Psychol.* **2001**, *86*, 114–121. [[CrossRef](#)] [[PubMed](#)]
36. Langer, E.J.; Blank, A.; Chanowitz, B. The mindlessness of ostensibly thoughtful action: The role of “placebic” information in interpersonal interaction. *J. Pers. Soc. Psychol.* **1978**, *36*, 635–642. [[CrossRef](#)]
37. Ahmad Husairi, M.; Rossi, P. Delegation of purchasing tasks to AI: The role of perceived choice and decision autonomy. *Decis. Support Syst.* **2024**, *179*, 114166. [[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.