iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** XI **Month of publication:** November 2026

DOI: https://doi.org/10.22214/ijraset.2026.85339

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue XI Nov 2026- Available at www.ijraset.com

Generative AI in Consumer Decision Making: How Algorithmic Recommendations Shape Brand Trust

Ms. Manisha Yogesh Thorat
Assistant professor, Department of Management Science, MIT-CSN, Maharashtra India

Abstract: ***Generative artificial intelligence (GenAI) is rapidly moving from the back end of recommendation engines to the consumer-facing interface, where conversational assistants, generated product descriptions and personalised feeds increasingly decide what shoppers see, compare and buy. While the literature on algorithm aversion and algorithm appreciation explains when consumers accept algorithmic advice, little is known about how such advice reshapes trust in the brand that deploys it. This paper develops and specifies a multi-level, testable framework in which four attributes of a GenAI recommendation (perceived personalisation, perceived accuracy and relevance, transparency and explainability, and disclosure of the AI source) influence brand trust through two opposing appraisal pathways: perceived algorithmic competence and fairness, which build trust, and perceived manipulation and privacy concern, which erode it. Brand trust in turn drives purchase intention and loyalty, with AI literacy and task type (utilitarian versus hedonic) as boundary conditions. Nine hypotheses are derived from trust theory, the persuasion knowledge model and the human–automation literature, and a dynamic trust-calibration loop is proposed in which each decision cycle updates the trust baseline used in the next. A two-study empirical design is specified, comprising a PLS-SEM survey and a 3 × 2 × 2 online experiment, and a Monte Carlo power analysis shows that a survey of about 300 respondents gives at least 95% power for the main hypothesised paths, while a 600-participant experiment gives about 92% power for medium-small effects of disclosure. The paper contributes a brand-level, trust-centred account of GenAI recommendations, identifies transparency as the lever that most plausibly converts algorithmic personalisation into trust, and derives managerial and regulatory implications. The study is conceptual with simulated power analysis; empirical estimates are not reported.***

Keywords: ***generative AI, algorithmic recommendation, brand trust, consumer decision making, algorithm aversion, transparency, explainability, persuasion knowledge, PLS-SEM.***

I. INTRODUCTION

Consumer decision making is increasingly mediated by algorithms. Search rankings, “you may also like” carousels, streaming queues and social feeds already decide which options consumers consider, and the arrival of generative AI (GenAI) adds a new layer: systems that converse, explain, compare and generate persuasive language and imagery in real time. A shopper can now describe a need in natural language and receive a tailored, fluent recommendation that appears to reflect understanding of their situation [25]. For brands, this promises scale, personalisation and lower service cost. It also changes the nature of what consumers are asked to trust. Trust is no longer placed only in a product, a retailer or a salesperson, but also in an opaque system that sits between the two [6], [13].

Brand trust is among the most consequential constructs in marketing. It reduces perceived risk, supports loyalty and underpins willingness to pay [8], [9], [10]. Classical models treat it as a function of the brand’s ability, benevolence and integrity as perceived by the consumer [7]. GenAI recommendations complicate this picture in two ways. First, a recommendation is a brand action: when an assistant steers a consumer towards a poor, biased or manipulative choice, the consumer may blame the brand and not the algorithm. Second, the same recommendation may be received either as evidence of competence and care (“it understands me”) or as evidence of surveillance and persuasion (“it is using what it knows about me to sell to me”) [21], [22]. Which reading prevails is an empirical and theoretical question that existing work has not resolved at the brand level.

Research on how consumers react to algorithms has grown quickly but has focused mainly on the acceptance of advice. Studies show that people sometimes discount algorithmic advice after seeing it err [1], sometimes prefer it to human advice [2], resist it in domains where they feel unique [3], and rely on it more for objective than for subjective tasks [4]. These findings concern the decision at hand. They say less about the downstream consequence that matters most to the firm, namely whether the brand deploying the algorithm is trusted more or less afterwards, and whether this effect accumulates over repeated interactions.

©IJRASET: All Rights are Reserved

This paper addresses that gap. It asks three research questions: (RQ1) Through which psychological mechanisms do the attributes of GenAI recommendations affect consumer trust in the brand? (RQ2) Which design choices, notably transparency and disclosure, make those mechanisms favourable and which make them harmful? (RQ3) How do consumer characteristics and task context moderate these effects, and how does trust evolve over repeated recommendation cycles?

The contributions are fourfold. First, the paper integrates trust theory, the persuasion knowledge model and the human–automation literature into a single brand-level framework with two opposing mediation paths. Second, it proposes nine falsifiable hypotheses and a dynamic trust-calibration loop. Third, it specifies a complete, pre-registrable two-study design and supports its sample-size choices with Monte Carlo and analytical power calculations. Fourth, it derives managerial and regulatory implications. The paper is conceptual and methodological; it does not report empirical estimates, and all numerical results presented are from simulations whose purpose is design validation.

The remainder of the paper is organised as follows. Section II reviews the literature. Section III develops the framework and hypotheses. Section IV presents the research design. Section V reports the power analysis. Section VI discusses the dynamic perspective. Section VII sets out theoretical and managerial implications, Section VIII covers ethical and regulatory considerations, and Sections IX and X give limitations, future research and conclusions.

II. THEORETICAL BACKGROUND AND LITERATURE REVIEW

A. GenAI and algorithmic recommendation in consumer decisions

Recommender systems have long been studied as decision aids that reduce search effort and improve choice quality [17]. Early work showed that characteristics of recommendation agents, such as the type of input they collect and the explanation they offer, shape both use and outcomes [17]. Marketing scholarship now frames artificial intelligence as a strategic capability that spans mechanical, thinking and feeling tasks, from data processing to personalised interaction [24], and as a force that changes how firms engage customers across the journey [23]. GenAI differs from earlier recommenders because it produces open-ended natural-language and multimodal content, simulates conversational competence, and can be prompted to adopt a persuasive style [25]. It therefore combines the informational role of a recommender with the social and persuasive role of a salesperson, which makes trust and persuasion knowledge central concerns.

Puntoni et al. provide an experiential account of how consumers meet AI, organised around four experiences: data capture, classification, delegation and social interaction. Each experience carries a characteristic tension, for example between feeling served and feeling exploited when data are captured, or between feeling empowered and feeling replaced when decisions are delegated [6]. This account is valuable because it shows why the same system can produce opposite emotional and evaluative responses, and it motivates the two opposing appraisal pathways proposed here.

B. Algorithm aversion and algorithm appreciation

Dietvorst, Simmons and Massey showed that people who watched an algorithm make errors became less willing to use it, even when it outperformed a human forecaster [1]. In contrast, Logg, Minson and Moore found that lay participants adhered more strongly to advice when it was described as coming from an algorithm than from a person, a pattern they termed algorithm appreciation [2].

Longoni, Bonezzi and Morewedge traced resistance to medical AI to “uniqueness neglect”, the belief that algorithms cannot account for the individual’s particular circumstances [3]. Castelo, Bos and Lehmann demonstrated that reliance depends on the task: consumers are more reluctant to rely on algorithms for subjective tasks, and anthropomorphic framing or evidence of objectivity can reduce this reluctance [4].

Yalcin and colleagues showed that consumers respond less favourably to positive decisions made by algorithms than to the same decisions made by humans, whereas responses to negative decisions are similar, which they linked to consumers’ attributions about their own deservingness [5].

Taken together, these studies suggest that the acceptance of algorithmic advice is conditional on task type, perceived uniqueness handling, prior experience with errors and the way outcomes are attributed. Sundar’s machine-heuristic account adds that people apply the rule “machines are objective and unbiased” in some contexts and that this heuristic can be strengthened or weakened by interface cues [14]. These results motivate task type and AI literacy as moderators in the present framework. However, none of these studies measures the effect on trust in the brand that provided the algorithm.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

C. Brand Trust

Brand trust is the willingness of the average consumer to rely on the ability of the brand to perform its stated function [8]. It is generally conceptualised as having cognitive elements (competence, reliability) and affective or relational elements (benevolence, integrity) [7], [9]. Morgan and Hunt placed trust, together with commitment, at the centre of relationship marketing, arguing that it mediates the effect of shared values and communication on cooperation and loyalty [9]. Erdem and Swait described brands as signals whose credibility reduces information costs and perceived risk, and credibility rests on trustworthiness and expertise [10]. Chaudhuri and Holbrook showed that brand trust and brand affect jointly drive purchase and attitudinal loyalty, with consequences for market share and price premium [8].

In online settings, trust additionally reduces the perceived risk of transacting with an unseen party and affects adoption of the technology through which transactions occur [15]. Komiak and Benbasat distinguished cognitive from emotional trust in recommendation agents and showed that personalisation and familiarity influence trust and intention to adopt [16]. This work is closest in spirit to the present paper but concentrates on trust in the agent, not in the brand behind it.

D. Trust in automation and in AI

The human–automation literature treats trust as an attitude that guides reliance when uncertainty and vulnerability are present, and stresses calibration: trust should match actual system capability, because over-trust produces misuse and under-trust produces disuse [11]. Hoff and Bashir synthesised the empirical evidence into three layers of trust variability: dispositional trust (stable traits), situational trust (context and task) and learned trust (experience with the specific system, which is itself shaped by initial expectations) [12]. Glikson and Woolley reviewed human trust in AI and distinguished cognitive from emotional trust, noting that tangibility, transparency, reliability and immediacy behaviours affect cognitive trust, while anthropomorphism and embodiment affect emotional trust [13]. Mayer et al.'s ability–benevolence–integrity model supplies the vocabulary for translating these ideas to the brand level [7].

E. Transparency, explainability and disclosure

Explainability is widely argued to be a precondition for trust in AI systems. Rai described a shift from "black box" to "glass box" models and discussed how explanations influence trust and acceptance [19]. Shin found that explainability and causability influence perceived transparency, fairness and trust, and thereby acceptance [18]. Explanations are not uniformly beneficial: they take effort to process and can reveal limitations; whether they help depends on whether they are comprehensible and match the user's level of expertise. Disclosure of the AI source is a distinct design decision. It informs the consumer that a machine, not a human, produced the advice, and thereby may activate persuasion knowledge, the lay theories people hold about marketers' tactics, which lead them to scrutinise and sometimes resist influence attempts [20]. The Elaboration Likelihood Model adds that consumers who process peripherally rely on cues such as source labels, whereas those who process centrally evaluate argument quality [26].

F. Personalisation, privacy and manipulation

Personalisation is a central promise of GenAI, but the personalisation–privacy paradox implies that the same personal fit that raises relevance can also raise privacy concern, especially when consumers do not know how information was collected [21]. Martin and Murphy argue that data privacy is a core marketing concern in which consent, control and perceived fairness of data use determine trust [22]. Aguirre et al. showed that personalised advertising can reduce effectiveness when it triggers vulnerability and that trust-building strategies mitigate this effect [21]. For GenAI, the concern extends from data use to persuasion: a system that can adapt tone, framing and emphasis to each individual may be seen as manipulative, even when its recommendations are accurate.

G. Research gap

The review reveals four gaps. (1) The dependent variable in most algorithm-reaction studies is acceptance of advice or choice, not brand-level trust. (2) Trust studies of recommendation agents treat the agent as the trustee, whereas consumers may hold the brand accountable. (3) Transparency, disclosure and personalisation are typically studied separately, although in deployed GenAI interfaces they are jointly determined design choices. (4) Static designs dominate, although trust is learned over repeated interactions [11], [12]. The framework below addresses these gaps.

III. CONCEPTUAL FRAMEWORK AND HYPOTHESES

Fig. 1 presents the framework. Four attributes of a GenAI recommendation, perceived personalisation (PP), perceived accuracy and relevance (PA), transparency and explainability (TX) and AI-source disclosure (D), are proposed to affect brand trust (BT) directly and through two appraisal mechanisms: perceived algorithmic competence and fairness (CF), which is trust-building, and perceived manipulation and privacy concern (PM), which is trust-eroding. BT then predicts purchase intention and brand loyalty (BL). AI literacy and task type moderate the paths. Table I defines the constructs.

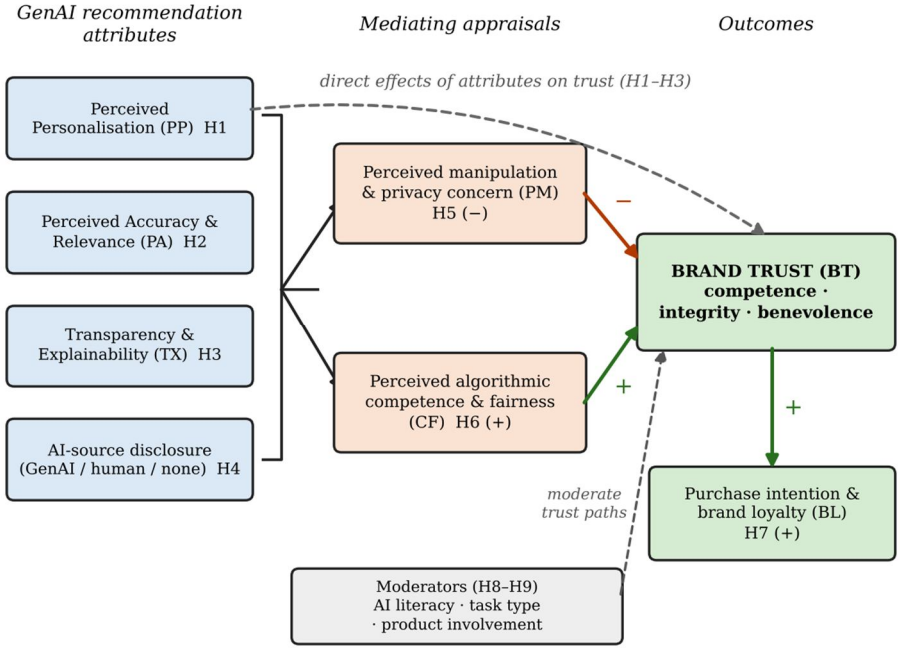


Fig. 1. Conceptual framework: GenAI recommendation attributes, dual appraisal mechanisms, brand trust and outcomes.

Table I. Construct Definitions and Theoretical Anchors

Construct	Definition (in this study)	Anchor
Perceived personalisation (PP)	Degree to which the recommendation is perceived to fit the consumer's own needs, context and preferences	[16], [21]
Perceived accuracy & relevance (PA)	Perceived quality and usefulness of the recommended options for the task at hand	[17]
Transparency & explainability (TX)	Perceived clarity of why a recommendation was made and what data and logic it rests on	[18], [19]
AI-source disclosure (D)	Explicit labelling that the recommendation is generated by AI (manipulated factor)	[20], [14]
Perceived manipulation & privacy concern (PM)	Belief that the system uses personal information or persuasive framing to influence the consumer against their interest	[20], [22]
Perceived algorithmic competence & fairness (CF)	Belief that the system is capable, unbiased and treats the consumer fairly	[7], [13]
Brand trust (BT)	Willingness to rely on the brand's ability, benevolence and integrity	[7], [8]
Brand loyalty & purchase intention (BL)	Intention to buy again, recommend and resist switching	[8], [27]
AI literacy	Knowledge of how AI systems work and their limits	[12]

A. Direct effects of recommendation attributes

H1. *Perceived personalisation positively affects brand trust.* A recommendation that fits the consumer signals that the brand understands and cares about the consumer's needs, which is evidence of both ability and benevolence [7], [16].

H2. *Perceived accuracy and relevance positively affects brand trust.* Accuracy is the most direct evidence of competence, and the human–automation literature consistently finds that observed reliability is the strongest driver of learned trust [11], [12].

H3. *Transparency and explain ability positively affects brand trust.* Explanations give the consumer a basis for judging the system's reasoning, which supports cognitive trust and perceived integrity [13], [18], [19].

B. Disclosure and the manipulation pathway

H4. *Disclosure of the AI source increases perceived manipulation and privacy concern when it is not accompanied by an explanation, and this effect is attenuated when an explanation is present.* Labelling alerts the consumer that a persuasion attempt is machine-generated and may activate persuasion knowledge [20]; an accompanying explanation gives the consumer a coping resource and makes the influence attempt legible. The interaction of disclosure and explanation is therefore the central experimental contrast.

H5. *Perceived manipulation and privacy concern negatively affects brand trust and mediates the effects of personalisation (positively, via the personalisation–privacy paradox [21]), transparency (negatively) and disclosure on trust.* Perceived manipulation conveys a lack of integrity and benevolence on the part of the brand [7], [22].

C. Competence and fairness pathway, and outcomes

H6. *Perceived algorithmic competence and fairness positively affects brand trust and mediates the effects of accuracy and transparency.* Consumers who judge the system as capable and fair attribute these qualities to the brand that deploys it, consistent with the transfer of trust from technology to its provider [15], [13].

H7. *Brand trust positively affects purchase intention and brand loyalty.* This replicates the established chain from trust to loyalty [8], [9] in the GenAI setting and offers a necessary check on the nomological network.

D. Boundary conditions

H8. *AI literacy strengthens the positive effect of transparency on brand trust and weakens the effect of perceived manipulation on trust.* Consumers who understand how AI systems work can use explanations and are less prone to over-attributing intent to the system [12].

H9. *Task type moderates the effect of accuracy: the effect of perceived accuracy on competence and trust is weaker for hedonic, taste-based decisions than for utilitarian decisions.* Consumers regard subjective tasks as less amenable to algorithmic judgement [4]; accuracy is therefore a less credible competence signal in such contexts and personalisation cues may matter more.

The structural relations can be written compactly. With standardised constructs and error terms ε, the mediators and the outcome are modelled as

$$PM = a_1 \cdot PP + a_2 \cdot TX + a_3 \cdot D + a_4 \cdot (D \times TX) + \varepsilon_1 \quad (1)$$

$$CF = b_1 \cdot PA + b_2 \cdot TX + b_3 \cdot PP + \varepsilon_2 \quad (2)$$

$$BT = c_1 \cdot PP + c_2 \cdot PA + c_3 \cdot TX + \gamma_1 \cdot PM + \gamma_2 \cdot CF + \varepsilon_3 \quad (3)$$

$$BL = \delta \cdot BT + \varepsilon_4 \quad (4)$$

The indirect effect of transparency on trust through manipulation concern is $a_2 \cdot \gamma_1$, and through competence and fairness it is $b_2 \cdot \gamma_2$. The hypotheses therefore predict $a_2 < 0$, $\gamma_1 < 0$, $b_2 > 0$ and $\gamma_2 > 0$, so that transparency has a positive total effect on trust that operates through two channels. The interaction term a_4 is expected to be negative, which reflects the buffering of disclosure by explanation.

IV. RESEARCH DESIGN

The framework is to be tested in two complementary studies. The survey (Study 1) assesses the structural model in naturalistic consumption; the experiment (Study 2) establishes causality for the manipulable design factors.

A. Study 1: Survey and PLS-SEM

Study 1 targets adult online shoppers who have used a GenAI shopping assistant or conversational recommender within the past three months, recruited through a panel provider with quota sampling on age, gender and region. A screening question ensures recent experience, and respondents recall a specific recent interaction with a named brand. The target sample is $N \geq 400$ after quality screening, which exceeds the 300 respondents suggested by the power analysis in Section V and allows for attrition and for multi-

group analysis. Partial least squares structural equation modelling (PLS-SEM) is the primary estimation method because the model is complex, includes composite constructs and interaction terms, and is prediction-oriented [30]; a covariance-based check will be reported as robustness.

B. Study 2: Online experiment

Study 2 is a between-subjects 3 (disclosure: none, "generated by AI", "curated by a human expert") × 2 (explanation: absent, present) × 2 (task: utilitarian, hedonic) experiment. Participants interact with a scripted GenAI shopping assistant for a fictitious but realistic brand, receive a recommendation with the assigned labelling and explanation, and then complete the measures. Recommendation quality is held constant across conditions through pre-tested content, so that effects cannot be attributed to differences in accuracy. Utilitarian tasks (for example choosing a laptop for work) and hedonic tasks (for example choosing a gift or a holiday experience) are pre-tested for perceived task type. A manipulation check, an attention check and a realism check are included. The target is N = 600, or about 50 participants per cell.

C. Measures

All latent constructs are measured with multi-item, seven-point Likert scales adapted from established instruments (Table II). Items will be pre-tested in a pilot (N ≈ 60) for clarity and reliability, and translated and back-translated where bilingual administration is needed. Control variables include age, gender, income, online shopping frequency, prior brand familiarity, dispositional trust and general attitude to AI.

TABLE II. MEASUREMENT PLAN

Construct	Operationalisation	Adapted from / guidance
Brand trust	Reflective, 3 dimensions: competence, integrity, benevolence	[7], [8]
Perceived manipulation	Reflective, 4–5 items on persuasion intent and unfair influence	[20]
Privacy concern	Reflective, 4 items on data use and control	[22]
Competence & fairness	Reflective, 5 items	[7], [13]
Personalisation	Reflective, 4 items on perceived fit	[16], [21]
Accuracy & relevance	Reflective, 4 items	[17]
Transparency & explain ability	Reflective, 4 items on clarity of reasoning	[18]
Purchase intention & loyalty	Reflective, 4 items	[8], [27]
AI literacy	Knowledge index and self-rating, 6 items	Developed for this study; see [12]

D. Analysis plan and validity safeguards

The measurement model will be assessed for indicator reliability, composite reliability, convergent validity using the average variance extracted [28], and discriminant validity using the Fornell–Larcker criterion and the heterotrait–monotrait ratio [29]. Common method bias will be limited by procedural remedies (anonymity, separation of predictor and criterion blocks, randomised item order) and tested statistically with a full collinearity assessment [30], [31]. The structural model will be estimated with bootstrapped confidence intervals for direct and indirect effects, and mediation will be assessed through the significance and sign of indirect effects [32]. Moderation will be tested with product-indicator or two-stage interaction terms, and measurement invariance will be examined before any multi-group comparison. In the experiment, analysis of variance and planned contrasts test H4, and a moderated mediation model tests H5, H6, H8 and H9. All hypotheses, sample sizes, exclusions and analyses will be pre-registered. Respondents will give informed consent, data will be anonymised, and the study will be submitted for institutional ethical review.

V. POWER ANALYSIS AND DESIGN VALIDATION

To confirm that the proposed samples are adequate, a Monte Carlo simulation of the structural model was run. Data were generated from the model in Eqs. (1) to (3) with moderate, theoretically plausible coefficients (competence and fairness → trust, β = 0.30; manipulation concern → trust, β = −0.25; transparency → manipulation concern, −0.35; accuracy → competence and fairness, 0.40; and small direct paths of 0.10), with correlated antecedents and a binary disclosure factor. For each sample size, 500 datasets were generated, and power was defined as the proportion of datasets in which the target path was significant at α = 0.05. Indirect effects were assessed by the joint-significance test, which is known to perform well in terms of power [32]. These are simulation results on artificial data and indicate only the precision of the design; they are not empirical findings.

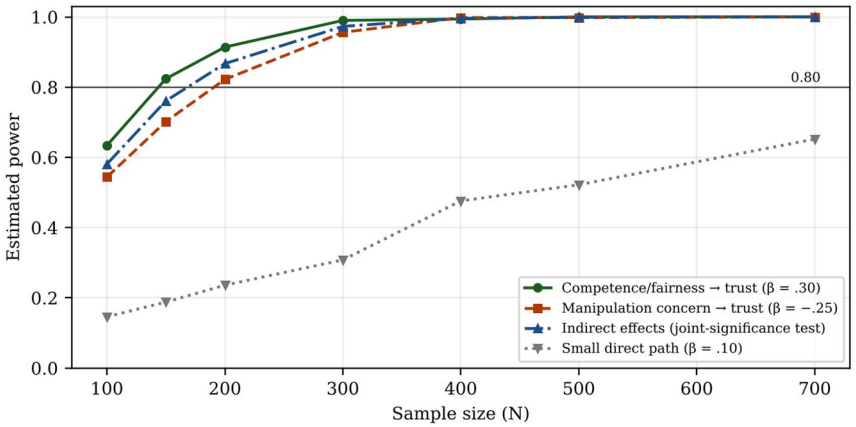


Fig. 2. Estimated power as a function of sample size for the main structural paths (500 simulated datasets per point).

Table III. Estimated Power By Sample SizE (MONTE CARLO, α = 0.05)

N	Competence/fairness → trust	Manipulation → trust	Indirect effects (mean)	Small direct path (β = .10)
100	63%	54%	58%	15%
150	82%	70%	76%	19%
200	91%	82%	87%	24%
300	99%	96%	97%	31%
400	99%	100%	100%	48%
500	100%	100%	100%	52%
700	100%	100%	100%	65%

The main hypothesised paths exceed 80% power at about 200 respondents and reach roughly 95% or more at 300. Small direct effects of 0.10 are not reliably detectable even at N = 700 (about 65% power), which has an important interpretive consequence: the study is powered to detect the mediated effects through CF and PM, but a non-significant direct effect of PP, PA or TX on trust should not be read as evidence of absence. The target of N ≥ 400 for Study 1 is therefore adequate for the central claims, and larger samples would be needed to establish small direct effects.

For Study 2, analytical power for the 12-cell design was computed from the non-central F distribution (Table IV). With N = 600, power exceeds 95% for a main effect of explanation of Cohen's f = 0.15 (about d = 0.30) and is about 92% for the three-level disclosure factor and for the disclosure × explanation interaction at f = 0.15, and above 100% at f = 0.20. Power for a small main effect of task type (f = 0.10) is lower (69%), but task type is primarily a moderator in this design and is tested through interaction contrasts.

TABLE IV. ANALYTICAL POWER FOR THE 3 × 2 × 2 EXPERIMENT (α = 0.05)

Total N	Explanation main (f = .15)	Disclosure main / interaction (f = .15)	Interaction (f = .20)	Task main (f = .10)
360	81%	72%	93%	47%
480	91%	84%	98%	59%
600	96%	92%	100%	69%

VI. A DYNAMIC VIEW: THE TRUST-CALIBRATION LOOP

Most studies of reactions to algorithms use a single interaction. In practice, consumers meet a GenAI recommender repeatedly, and trust at time t is both an outcome of past interactions and an input to how the next one is interpreted [11], [12]. Fig. 3 shows the proposed loop. The brand deploys a recommender; the consumer receives a recommendation together with whatever disclosure and explanation the interface provides; the consumer appraises its relevance, fairness and manipulativeness; brand trust is updated; behaviour follows (click, accept, buy, return or switch); and the observed outcome (the recommendation proved good, poor or deceptive) feeds back into the next appraisal.

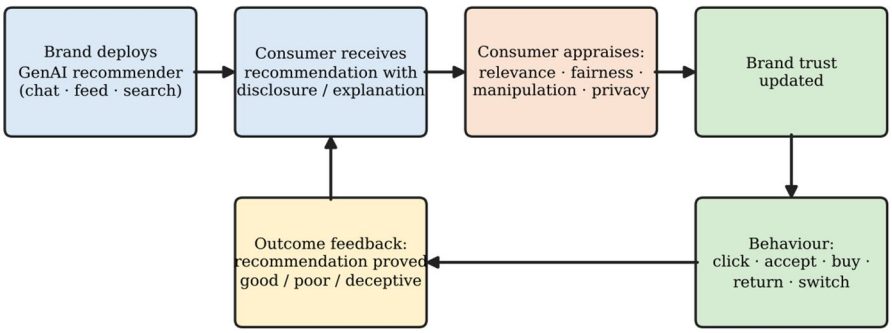


Fig. 3. Dynamic trust-calibration loop between consumer, GenAI recommender and brand.

The loop yields propositions that complement the nine hypotheses and are suited to longitudinal or diary designs. A simple updating rule illustrates the logic. Let T(t) be trust at cycle t and o(t) ∈ [0, 1] the experienced outcome quality; then

$$T(t+1) = T(t) + \lambda(t)\cdot[\, o(t) - E(t)\,] \qquad (5)$$

where E(t) is the expectation the consumer formed before the interaction and λ(t) is a learning rate. This form captures calibration [11]: trust rises when outcomes exceed expectations and falls when they fall short, and the size of the update depends on λ(t), which is expected to be larger after a salient failure, consistent with the finding that observing errors weighs heavily on subsequent use [1].

P1. *Asymmetry.* Trust losses after a perceived manipulative or deceptive recommendation are larger and longer-lasting than gains after an equally good one (λ is larger for negative than for positive discrepancies).

P2. *Expectation setting.* Disclosure and explanation influence E(t), so that transparent interfaces reduce the size of negative discrepancies and thereby protect trust when errors occur.

P3. *Spillover.* Trust updates generalise from the recommender to the brand when the brand is clearly identified as the provider and weaken when a third-party platform is salient.

These propositions are not tested in the two studies above but define the agenda for a longitudinal extension.

VII. DISCUSSION AND IMPLICATIONS

A. *Theoretical implications*

First, the framework moves the focus from the acceptance of advice to the trust that results for the brand, and thereby links the algorithm-reaction literature [1]–[5] with the brand-trust tradition [7]–[10].

Second, it specifies two opposing mediating mechanisms, which explains why GenAI recommendations can both build and damage trust, and accounts for the mixed evidence of aversion and appreciation. Third, by treating disclosure and explanation jointly, it brings together the persuasion knowledge and explainable-AI literatures, which have so far developed largely independently [19], [20]. Fourth, the dynamic loop positions brand trust as a calibrated state variable in the sense of the automation literature [11], not a fixed attitude.

B. Managerial implications

Table V summarises design levers and the mechanisms and risks associated with each. The central managerial message is that personalisation, the most visible benefit of GenAI, is double-edged and its trust payoff depends on transparency. A brand that personalises without explaining invites the manipulation reading; a brand that explains what the recommendation is based on and gives the consumer control converts personal fit into evidence of benevolence.

TABLE V. DESIGN LEVERS, MECHANISMS AND RISKS

Design lever	Expected mechanism	Principal risk
Clear "why this recommendation" explanation	Raises TX and CF; lowers PM	Information overload; revealing limits
Disclosing the AI source	Honest signalling; supports integrity	Persuasion-knowledge activation if unexplained
Data-use controls and consent	Reduces privacy concern	Friction; lower personalisation depth
Calibrated confidence statements	Sets E(t); protects trust after errors	Appearing uncertain
Human escalation path	Addresses uniqueness concerns	Cost
Task-sensitive design (utilitarian vs hedonic)	Matches cues to what is credible	Over-segmentation

Managers should also monitor trust as a leading indicator of the health of recommendation systems, instead of relying only on click-through or conversion, which can rise in the short term while trust erodes. Finally, the finding that small direct effects require large samples suggests that A/B tests of trust-related design changes need to be adequately powered, or they will systematically miss modest but meaningful effects.

VIII. ETHICAL AND REGULATORY CONSIDERATIONS

The line between helpful personalisation and manipulation is also a regulatory one. The EU Artificial Intelligence Act introduces transparency obligations for AI systems that interact with people and for AI-generated content, and prohibits certain manipulative practices [33]. In India, the Digital Personal Data Protection Act, 2023 sets out consent and purpose-limitation requirements for the processing of personal data [34], which directly affects the data capture that personalised GenAI depends on. These regimes are consistent with the framework's prediction that transparency and control are conditions of trust: designs that meet them are likely to be both compliant and more trusted. Researchers and practitioners must also guard against using the findings to craft more effective covert persuasion. The experimental stimuli in Study 2 should be debriefed, and the research should be framed around making AI-mediated influence legible to consumers. Fairness is a further concern: personalisation that systematically disadvantages groups would harm perceived fairness and could invite legal liability.

IX. LIMITATIONS AND FUTURE RESEARCH

The paper has several limitations. It is conceptual and the empirical claims are hypotheses; the numerical results are simulations and show design adequacy, not effect sizes. The model is simplified: trust is treated as one construct with three facets, and affective trust and brand affect [8] are not modelled separately. The survey relies on self-reports and recall of past interactions, which are subject to bias, and the experiment uses a scripted assistant that is less rich than a live GenAI system.

The context is online retail, and results may not transfer to high-stakes domains such as finance or health, where the consequences of error and the regulatory setting differ. Future research should (i) run the two studies and report effect sizes with intervals; (ii) test the dynamic loop with longitudinal or diary designs and field data; (iii) examine cross-cultural differences, particularly between emerging and advanced markets; (iv) study the role of anthropomorphic and voice-based interfaces; (v) investigate multi-brand ecosystems in which a platform's assistant recommends other brands' products, which raises the question of whom the consumer holds accountable; and (vi) examine the effects of hallucinated or fabricated content, which is a failure mode specific to generative systems.

X. CONCLUSION

GenAI is changing consumer decision making by placing a persuasive, conversational system between brands and their customers. This paper argued that the effect on brand trust is not uniformly positive or negative but is determined by two competing appraisals: whether the recommendation is read as the work of a competent and fair system, or as an exercise in manipulation and data exploitation. Transparency and explainability are the levers most likely to tilt this balance, and disclosure of AI involvement is safest when accompanied by explanation. The paper offered nine hypotheses, a dynamic trust-calibration view, a complete two-study design with validated sample sizes, and a set of managerial and regulatory implications. Empirical testing of the framework is the immediate next step. If the hypotheses hold, brands will be able to treat transparency not as a compliance cost but as a source of lasting brand trust.

XI. ACKNOWLEDGMENT

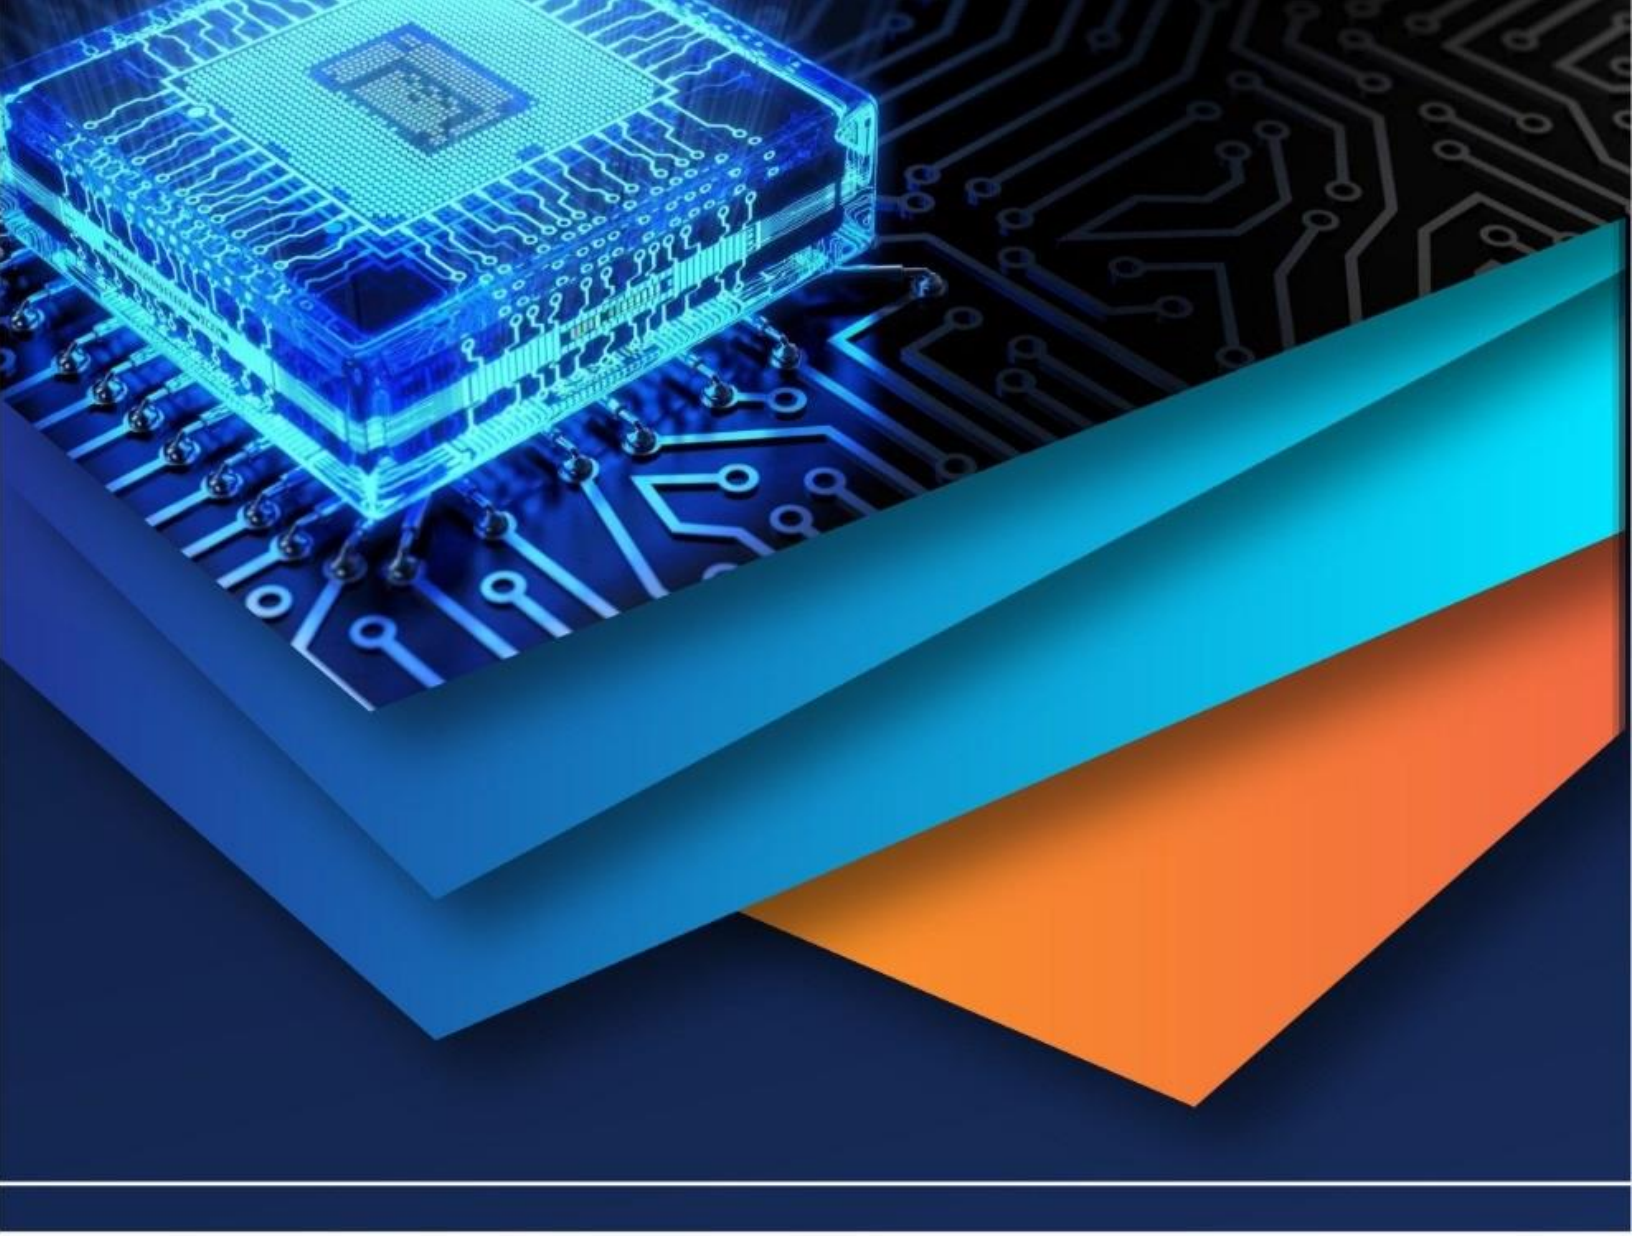
The author thanks colleagues and the institution for their support.

REFERENCES

[1] B. J. Dietvorst, J. P. Simmons and C. Massey, "Algorithm aversion: People erroneously avoid algorithms after seeing them err," Journal of Experimental Psychology: General, vol. 144, no. 1, pp. 114–126, 2015.
[2] J. M. Logg, J. A. Minson and D. A. Moore, "Algorithm appreciation: People prefer algorithmic to human judgment," Organizational Behavior and Human Decision Processes, vol. 151, pp. 90–103, 2019.
[3] C. Longoni, A. Bonezzi and C. K. Morewedge, "Resistance to medical artificial intelligence," Journal of Consumer Research, vol. 46, no. 4, pp. 629–650, 2019.
[4] N. Castelo, M. W. Bos and D. R. Lehmann, "Task-dependent algorithm aversion," Journal of Marketing Research, vol. 56, no. 5, pp. 809–825, 2019.
[5] G. Yalcin, S. Lim, S. Puntoni and S. M. J. van Osselaer, "Thumbs up or down: Consumer reactions to decisions by algorithms versus humans," Journal of Marketing Research, vol. 59, no. 4, pp. 696–717, 2022.
[6] S. Puntoni, R. W. Reczek, M. Giesler and S. Botti, "Consumers and artificial intelligence: An experiential perspective," Journal of Marketing, vol. 85, no. 1, pp. 131–151, 2021.
[7] R. C. Mayer, J. H. Davis and F. D. Schoorman, "An integrative model of organizational trust," Academy of Management Review, vol. 20, no. 3, pp. 709–734, 1995.
[8] A. Chaudhuri and M. B. Holbrook, "The chain of effects from brand trust and brand affect to brand performance: The role of brand loyalty," Journal of Marketing, vol. 65, no. 2, pp. 81–93, 2001.
[9] R. M. Morgan and S. D. Hunt, "The commitment–trust theory of relationship marketing," Journal of Marketing, vol. 58, no. 3, pp. 20–38, 1994.
[10] T. Erdem and J. Swait, "Brand equity as a signaling phenomenon," Journal of Consumer Psychology, vol. 7, no. 2, pp. 131–157, 1998.
[11] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," Human Factors, vol. 46, no. 1, pp. 50–80, 2004.
[12] K. A. Hoff and M. Bashir, "Trust in automation: Integrating empirical evidence on factors that influence trust," Human Factors, vol. 57, no. 3, pp. 407–434, 2015.
[13] E. Glikson and A. W. Woolley, "Human trust in artificial intelligence: Review of empirical research," Academy of Management Annals, vol. 14, no. 2, pp. 627–660, 2020.
[14] S. S. Sundar, "Rise of machine agency: A framework for studying the psychology of human–AI interaction (HAII)," Journal of Computer-Mediated Communication, vol. 25, no. 1, pp. 74–88, 2020.
[15] D. Gefen, E. Karahanna and D. W. Straub, "Trust and TAM in online shopping: An integrated model," MIS Quarterly, vol. 27, no. 1, pp. 51–90, 2003.
[16] [16] S. Y. Komiak and I. Benbasat, "The effects of personalization and familiarity on trust and adoption of recommendation agents," MIS Quarterly, vol. 30, no. 4, pp. 941–960, 2006.
[17] B. Xiao and I. Benbasat, "E-commerce product recommendation agents: Use, characteristics, and impact," MIS Quarterly, vol. 31, no. 1, pp. 137–209, 2007.
[18] D. Shin, "The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI," International Journal of Human-Computer Studies, vol. 146, art. 102551, 2021.
[19] A. Rai, "Explainable AI: From black box to glass box," Journal of the Academy of Marketing Science, vol. 48, no. 1, pp. 137–141, 2020.
[20] M. Friestad and P. Wright, "The persuasion knowledge model: How people cope with persuasion attempts," Journal of Consumer Research, vol. 21, no. 1, pp. 1–31, 1994.
[21] C. Aguirre, D. Mahr, D. Grewal, K. de Ruyter and M. Wetzels, "Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness," Journal of Retailing, vol. 91, no. 1, pp. 34–49, 2015.
[22] K. D. Martin and P. E. Murphy, "The role of data privacy in marketing," Journal of the Academy of Marketing Science, vol. 45, no. 2, pp. 135–155, 2017.

[23] T. H. Davenport, A. Guha, D. Grewal and T. Bressgott, "How artificial intelligence will change the future of marketing," Journal of the Academy of Marketing Science, vol. 48, no. 1, pp. 24–42, 2020.
[24] M.-H. Huang and R. T. Rust, "A strategic framework for artificial intelligence in marketing," Journal of the Academy of Marketing Science, vol. 49, no. 1, pp. 30–50, 2021.
[25] Y. K. Dwivedi et al., "Opinion paper: 'So what if ChatGPT wrote it?' Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy," International Journal of Information Management, vol. 71, art. 102642, 2023.
[26] R. E. Petty and J. T. Cacioppo, Communication and Persuasion: Central and Peripheral Routes to Attitude Change. New York, NY, USA: Springer, 1986.
[27] I. Ajzen, "The theory of planned behavior," Organizational Behavior and Human Decision Processes, vol. 50, no. 2, pp. 179–211, 1991.
[28] C. Fornell and D. F. Larcker, "Evaluating structural equation models with unobservable variables and measurement error," Journal of Marketing Research, vol. 18, no. 1, pp. 39–50, 1981.
[29] J. Henseler, C. M. Ringle and M. Sarstedt, "A new criterion for assessing discriminant validity in variance-based structural equation modeling," Journal of the Academy of Marketing Science, vol. 43, no. 1, pp. 115–135, 2015.
[30] J. F. Hair, G. T. M. Hult, C. M. Ringle and M. Sarstedt, A Primer on Partial Least Squares Structural Equation Modeling (PLS-SEM), 3rd ed. Thousand Oaks, CA, USA: Sage, 2022.
[31] P. M. Podsakoff, S. B. MacKenzie, J.-Y. Lee and N. P. Podsakoff, "Common method biases in behavioral research: A critical review of the literature and recommended remedies," Journal of Applied Psychology, vol. 88, no. 5, pp. 879–903, 2003.
[32] D. P. MacKinnon, C. M. Lockwood, J. M. Hoffman, S. G. West and V. Sheets, "A comparison of methods to test mediation and other intervening variable effects," Psychological Methods, vol. 7, no. 1, pp. 83–104, 2002.
[33] European Parliament and Council, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)," Official Journal of the European Union, 2024.
[34] Government of India, "The Digital Personal Data Protection Act, 2023," Act No. 22 of 2023, 2023.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)