

From Feed to Conversation: How AI Brand Persona Continuity Shapes Social Presence and Consumer Trust

Rohan Kumar HM

rohan.kumarhm25pb@ibsindia.org

ICFAI Foundation for Higher Education

Dr. Clement Sudhahar

ICFAI Foundation for Higher Education

Dr. Sumanjit Das

ICFAI Foundation for Higher Education

Article

Keywords: AI brand persona, persona continuity, chatbot trust, social presence, parasocial familiarity, virtual influencer

Posted Date: August 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10413097/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Brands increasingly rely on artificial intelligence to maintain recognizable brand voice across social media, product communication, and chatbot-based service encounters. This conceptual paper introduces the Persona-Avatar Continuity-Trust (PACT) Model to explain how consumers' familiarity with an AI persona in social media may shape trust when they later interact with a recognizably continuous chatbot avatar. Drawing on parasocial relationship theory, virtual influencer research, anthropomorphism, social presence, dual-route trust formation, AI disclosure, expectation-confirmation reasoning, self-determination theory, self-brand congruity, and cross-cultural marketing, the paper develops an integrative theory-building framework. The model proposes that threshold-level familiarity with a persona can pre-activate social presence in later chatbot interactions, enabling a persona carryover effect that supports affective and cognitive trust. This process is conditional on avatar-persona continuity, truthful disclosure, perceived competence, performance-expectation fit, consumer digital literacy, persona selection agency, machine-heuristic preferences, and cultural context. The paper advances eight testable propositions linking persona continuity to trust formation and downstream outcomes such as engagement, loyalty, and brand equity. It shifts the unit of analysis from isolated AI touchpoints to continuous AI-mediated brand relationships across content and conversation. The framework also offers practical guidance for managing persona drift, AI disclosure, human escalation, data protection, and meaningful consumer choice. Future research should validate the model through staged experiments, longitudinal exposure designs, measurement development, non-avatar controls, and cross-cultural field studies.

1. INTRODUCTION

Consumers might repeatedly encounter an AI brand persona before they even have a service need. For instance, an Indian beauty brand might employ an expressive digital character to share ingredient narratives on Instagram, address common inquiries, and feature in product education videos. However, related virtual-idol research also suggests cross-cultural variation, particularly in East Asian contexts where virtual personae may become linked to mainstream or national identity formation (Song et al., 2024). When the same character reappears in a chatbot interface, the user may not perceive the interaction as a cold start with an unfamiliar system. Instead, the interaction can be viewed as an extension of an existing relationship that is well-known but inherently unequal. This is crucial because AI brand characters are increasingly assuming multiple organizational roles. While virtual influencers and brand avatars are employed for content creation, storytelling, and endorsements, conversational agents are utilized for providing recommendations and services. Still, theoretical traditions dealing with these functions have typically examined them in isolation. Research on parasocial relationships, social media influencers, and virtual influencers examines factors such as familiarity, authenticity, disclosure, engagement, and parasocial response within content environments (Hoffner & Bond, 2022; Lou, 2022; Lou et al., 2023; Stein et al., 2024). Research on chatbots more frequently examines trust established within bounded dialogues through dimensions such as warmth, competence, empathy, disclosure, and perceived reliability (Kull et al., 2021; Cheng et al., 2022; Huang et al., 2022; Prakash & Das, 2024; Wang

et al., 2026). This separation has significant consequences. Evidence suggests that anthropomorphic and social cues can shape trust-related responses in AI-mediated interactions, while disclosure of non-human identities may yield different outcomes depending on whether the context involves customer-service interactions or digital endorsements (De Cicco et al., 2021; Mozafari et al., 2022; Wang & Qiu, 2024). Still, these studies do not determine whether prior familiarity with the same recognizable AI persona alters the initial conditions of chatbot trust formation.

A consumer familiar with a digital persona may approach the interaction with expectations of warmth, tone, competence, and authenticity that are absent in encounters with an unfamiliar chatbot. This study addresses this theoretical gap by introducing the Persona-Avatar Continuity-Trust (PACT) Model. In the model, persona refers to the content-phase identity through which a brand's AI character becomes familiar, and avatar refers to the conversational instantiation through which that character responds in an interactive setting. Continuity refers to perceived consistency in visual cues, linguistic style, values, disclosure, and expected conduct across contexts. The central claim is intentionally provisional: persona-specific parasocial familiarity may pre-activate perceived social presence in a later chatbot encounter, thereby altering the route and rate at which trust is formed. This proposed mechanism is termed the persona carryover effect, referring to the transfer of relational warmth and readiness from asynchronous content exposure to synchronous or quasi-synchronous interactions within the same consumer-persona relationship. The effect remains conditional on a temporal trade-off: functional competence, including accuracy and safety, is a prerequisite for initial trust before relational warmth can become dominant. The machine heuristic may also moderate the process, because some users prefer nonanthropomorphic interfaces that they perceive as more objective, systematic, or efficient (Sundar & Kim, 2019; Yang & Sundar, 2024). The model culminates in eight propositions, including the final claim that affective and cognitive trust serve as primary drivers of brand engagement, loyalty, and equity. Three conditions make the proposed effect falsifiable: familiarity must reach a meaningful threshold; transfer should add to, rather than replace, competent within-chat interaction; and it should be specific to familiarity with the persona rather than general awareness of the brand. PACT also identifies a negative counterpart. When a consumer recognizes that the chatbot is intended to be a familiar persona but experiences its style, values, disclosure, or competence as inconsistent, the prior relationship may increase rather than reduce risk. This study refers to this as the continuity violation effect. The model further considers whether consumer persona selection agency, disclosure architecture, performance expectation fit, AI aversion, digital literacy, and cultural context alter these relationships. The study is guided by three research questions:

(RQ1) Under what conditions might parasocial familiarity with an AI brand persona transfer into a later chatbot encounter?

(RQ2) Through which social presence and cognitive-affective trust pathways might such transfers influence brand-related outcomes?

(RQ3) When might continuity, disclosure, performance, and consumer choice strengthen, weaken, or reverse the proposed effects?

The manuscript proceeds by stating its conceptual development approach, describing the data basis of the conceptual synthesis, reviewing the theoretical foundations, presenting the four-stage model and propositions, and specifying an empirical validation agenda and governance implications.

2. CONCEPTUAL DEVELOPMENT APPROACH AND STATUS OF CLAIMS

This article is written as an integrative conceptual paper rather than a systematic literature review or empirical study. Its purpose is to unify previously isolated theoretical perspectives around a specific unit of analysis: the recognizable AI brand persona that evolves from feed-based content to conversational services. This study does not purport to provide a comprehensive overview of all AI marketing research, nor does it report observed effect sizes. It employs published evidence to formulate plausible premises and generates new propositions that require further testing. The conceptual synthesis revolves around six key questions: whether audiences can develop parasocial familiarity with artificial or virtual figures; how anthropomorphic, authentic, and brand-value cues foster such familiarity; how social presence and trust operate within conversational AI; why disclosure effects vary across contexts; how continuity can create expectation confirmation or violation; and how consumer choice may enhance trust through autonomy and self-congruity. While classical theories offer construct definitions, newer peer-reviewed studies provide empirical anchors for AI and virtual influencer contexts. To enhance theoretical parsimony, established empirical premises are treated as foundational rather than being listed separately. The eight propositions below synthesize the paper's core theoretical contributions: content-to-conversation persona carryover, continuity violation, dual-route trust conversion, disclosure architecture, performance-expectation fit, persona selection agency, digital-literacy boundary conditions, and the influence of affective and cognitive trust on brand outcomes.

Following guidance on conceptual-article design, the paper is positioned as a hybrid of theory synthesis and model-building: it integrates constructs from adjacent literatures, specifies a new cross-touchpoint mechanism, and translates that mechanism into propositions that can be empirically challenged (MacInnis, 2011; Gilson & Goldberg, 2015; Jaakkola, 2020). The value of the article is therefore assessed in terms of theoretical clarity, internal coherence, novelty, and capacity to guide future empirical work rather than statistical representativeness.

By design, the model prioritizes content by investigating what occurs when social media familiarity precedes chatbot interaction. While a chatbot-first approach is a plausible option included in the research agenda, it is not assumed to be symmetrical. Similarly, the proposed Indian validation context is justified as a substantive empirical setting for digital brand relationships rather than as proof that the propositions have already been established in India. This positioning clarifies the scope of the conceptual claim: PACT advances testable relationships to be validated rather than presenting them as

established empirical effects. The theoretical foundations and their role in the PACT model are summarized in Table 1.

Table 1
Theoretical foundations and their role in the PACT model

Theory/literature	Established focus	Role in PACT	Claim status
Parasocial and transparasocial relations	Familiarity, perceived intimacy, and one-sided or asymmetric connection with mediated figures	Defines persona-specific familiarity accumulated through repeated feed-based content exposure	Empirical anchor
Virtual influencer authenticity research	Human-likeness, narrative coherence, transparency, source credibility, and brand outcomes in virtual influencer contexts	Grounds the conditions under which an AI brand persona can inherit authenticity, credibility, and brand-value associations from prior content exposure	Empirical anchor
Anthropomorphism and social presence	Humanlike cues, perceived agency, perceived humanness, and sensed interpersonal presence	Explains how prior exposure may pre-activate social presence before chatbot interaction, while also creating calibration risks if the conversational persona fails to match expectations	Anchor + extension
Machine heuristic	Preference for computer objectivity cues, procedural neutrality, systematic processing, and perceptions that machines are less socially biased	Bounds the relational logic by explaining why some consumers may trust machine-like neutrality more than anthropomorphic continuity or relational warmth	Boundary condition / moderator
Cognitive and affective trust	Competence-based trust and relational or emotion-based trust	Differentiates the mechanisms through which a familiar AI persona may convert content-based affect into service-context trust	Anchor + extension
Disclosure and persuasion knowledge	Identity cues, AI transparency, sponsorship recognition, skepticism, and consumer persuasion inferences	Specifies how timing, clarity, and context of AI disclosure may shape whether continuity is experienced as trustworthy transparency or as persuasive manipulation	Synthesis + extension
Expectation-confirmation and	Expectation confirmation, disconfirmation,	Explains why persona continuity can strengthen	Extension

Note. 'Empirical anchor' indicates that prior research provides evidence for the construct or relationship used by the PACT model. 'Anchor + extension' indicates that established theory supports the construct, while PACT extends it to the content-to-conversation transition of a recognizable AI brand persona. 'Extension' identifies a testable claim advanced by PACT. Evidence on persona stability and drift should be treated as emerging rather than established, especially when drawing on preprint work such as Brandt (2025) or Brandt and Wang (2025). Product involvement is treated as a distinct moderator rather than being folded into cultural theory or disclosure architecture. The machine heuristic is included as a boundary condition because some consumers may interpret machine-like neutrality as more trustworthy than relational warmth.

Theory/literature	Established focus	Role in PACT	Claim status
expectancy-violation logic	satisfaction formation, consistency, and violation of prior expectations	trust when the chatbot confirms prior expectations, but weaken trust when it produces mismatch, inconsistency, or persona drift	
Product involvement	Consumer motivation, perceived risk, and decision relevance in high- versus low-involvement product contexts	Acts as a boundary condition shaping whether AI persona cues, disclosure, and endorsement style are processed superficially or more critically	Empirical anchor + moderator
Self-determination and self-congruity	Autonomy, relatedness, identity fit, and perceived choice	Explains why consumer ability to select or align with a preferred AI persona may enhance trust through autonomy and self-congruity	Extension
Cultural theory and cross-cultural marketing	Cultural variation in trust, authenticity, identity, technology acceptance, and communication norms	Moderates how consumers interpret AI authenticity, disclosure, social presence, and brand-persona continuity across markets	Extension with empirical support
<i>Note. 'Empirical anchor' indicates that prior research provides evidence for the construct or relationship used by the PACT model. 'Anchor + extension' indicates that established theory supports the construct, while PACT extends it to the content-to-conversation transition of a recognizable AI brand persona. 'Extension' identifies a testable claim advanced by PACT. Evidence on persona stability and drift should be treated as emerging rather than established, especially when drawing on preprint work such as Brandt (2025) or Brandt and Wang (2025). Product involvement is treated as a distinct moderator rather than being folded into cultural theory or disclosure architecture. The machine heuristic is included as a boundary condition because some consumers may interpret machine-like neutrality as more trustworthy than relational warmth.</i>			

3. DATA

This article is conceptual and theory-building in nature. No primary empirical dataset was collected, and no human participants, customer records, platform logs, private organizational data, or experimental responses were used to generate the manuscript's propositions. The data basis of the study is therefore a secondary conceptual source corpus rather than a newly created empirical sample.

The source corpus consists of the cited works listed in the References section, including peer-reviewed journal articles, foundational theory papers, conference proceedings, scholarly books, arXiv preprints used only as emerging technical evidence, and one publicly available industry report on the Indian internet context. These sources were selected because they directly inform the theoretical traditions used in the PACT model: parasocial relationship theory, virtual influencer research, anthropomorphism and social presence, dual-route trust formation, disclosure and persuasion knowledge, expectation-

confirmation and expectancy-violation logic, self-determination theory, self-brand congruity, digital literacy, and cross-cultural marketing.

The data collection procedure followed an integrative conceptual synthesis process rather than a systematic review protocol. Relevant sources were identified through targeted literature searching, backward and forward citation chaining, and screening for conceptual relevance to AI brand personae, virtual influencers, chatbots, disclosure, social presence, trust, and consumer-brand relationship outcomes. Sources were retained when they supplied either a foundational theoretical construct, empirical evidence for a mechanism used in the model, or a boundary condition relevant to content-to-conversation continuity. Because the objective is theory construction rather than comprehensive evidence mapping, the process was not designed as a scoping review, did not use PRISMA-ScR procedures, and does not claim exhaustive coverage of every empirical study in AI marketing.

The analytical procedure involved qualitative concept mapping and theory synthesis. Constructs were compared across the selected literatures, overlapping mechanisms were consolidated, and relationships were organized into the four-stage Persona-Avatar Continuity-Trust architecture. The propositions were developed by linking the accumulated source evidence to theoretically plausible but as-yet untested relationships. No statistical analysis, machine-learning model training, scraping, survey analysis, or proprietary analytics platform was used in the study.

Data accessibility is limited to the cited sources themselves. Public reports and open-access preprints can be accessed from their issuing organizations or repositories, while journal articles and books may require publisher access or institutional subscription. No third-party permissions, data-use agreements, or ethical approvals were required because the manuscript does not reproduce restricted datasets, collect personal data, or analyze identifiable human-subject information. Any future empirical validation proposed in the manuscript should obtain informed consent, proportionate data protection safeguards, platform permissions where applicable, and institutional ethics approval whenever human participants or behavioral data are involved.

4. THEORETICAL FOUNDATIONS

4.1 Parasocial familiarity and mediated AI personae

Horton and Wohl (1956) defined parasocial interaction as a sense of intimacy experienced at a distance from media figures. Subsequent parasocial research shows that such perceived relationships can become meaningful features of media use without becoming reciprocal interpersonal relationships (Rubin et al., 1985; Hoffner & Bond, 2022). While contemporary digital platforms offer features such as comments and reactions that suggest responsiveness, an asymmetry often persists, allowing consumers to feel closeness without reciprocal social awareness. Lou (2022) characterizes this dynamic as a trans-parasocial relation within influencer contexts. For an AI brand persona, this implies that consumers distinguish software from humans and do not equate the relationship with interpersonal

friendship. Instead, repeated exposure to a consistent character can foster familiarity, an expected tone, and a feeling of relational warmth. Research comparing virtual and human influencers suggests that consumers can develop parasocial responses to virtual figures, particularly when they perceive similarity, human-likeness, authenticity, or credible persona cues (De Cicco et al., 2024; Lou et al., 2023; Stein et al., 2024). PACT therefore treats parasocial familiarity as a potential antecedent relational resource rather than simply as an outcome of content engagement.

4.2 Persona authenticity, brand-value inheritance, and anthropomorphism calibration

Virtual personas are intentionally crafted rather than naturally emergent brand representatives. Their perceived authenticity may therefore depend less on claims of human origin and more on the coherence of their narrative, voice, values, and transparent admission of artificiality. Allal-Chérif et al. (2024) suggest that the effectiveness of AI-powered virtual influencers cannot be reduced to appearance alone, because narrative and perceived authenticity are central to their marketing performance. Similarly, Lou et al. (2023) show that consumers may perceive virtual influencers as “authentically fake,” accepting their staged or artificial nature when the persona is narratively coherent and transparently constructed. Recent work further links anthropomorphised virtual influencer cues and perceived authenticity to consumer trust and engagement, including both cognitive and affective trust-related outcomes (Goel & Garg, 2025). PACT introduces brand-value inheritance to capture consumers’ perception that an AI persona recognizably embodies the brand’s values, personality, and narrative. This is a consumer-perceived construct rather than a claim about technical training data. Brand-value inheritance may foster parasocial familiarity because a consistent persona is easier to recognize, interpret, and remember across content and service contexts. Anthropomorphism calibration is treated as a distinct design condition: humanlike cues can enhance social legibility, but excessive realism or overly humanlike expectations may make deficiencies more salient and create discomfort or trust risk (Epley et al., 2007; Mori et al., 2012; Uysal et al., 2022; Gutuleac et al., 2024; Dabiran et al., 2024).

4.3 Social presence and dual-route trust formation in chatbot interaction

Social presence refers to the extent to which a mediated interaction is experienced as socially present or personally involving. In digital service contexts, social presence has been linked to consumer trust, and in chatbot contexts, responsiveness, conversational warmth, competence, empathy, reliability, and relational cues may contribute to trust-related evaluation (Gefen & Straub, 2004; Nordheim et al., 2019; Kull et al., 2021; Cheng et al., 2022; Huang et al., 2022; Prakash & Das, 2024; Wang et al., 2026). Trust should not be treated as a unitary judgement in this context. Cognitive trust concerns competence, reliability, accuracy, and appropriate oversight, whereas affective trust concerns perceived benevolence, warmth, emotional comfort, and relational confidence (McAllister, 1995; McKnight & Chervany, 2001; Dunn & Schweitzer, 2005; Johnson & Grayson, 2005). The PACT extension suggests that parasocial familiarity may alter the initial balance between these trust routes. An unfamiliar chatbot may initially be

judged mainly through competence and reliability cues; a recognized, familiar AI persona may make affective trust cues salient earlier, provided that the system performs adequately. This distinction matters for PACT because persona continuity is expected to change the sequencing of trust formation, not to replace functional competence as a requirement for trustworthy service.

4.4 Disclosure architecture and persuasion knowledge

Disclosure

evidence is not consistently positive or negative. In customer-interaction contexts, Luo et al. (2019) found that revealing a chatbot's identity before a sales conversation can decrease purchases by lowering perceived knowledge and empathy. De Cicco et al. (2025) similarly show that non-human identity cues can shape consumer responses through social presence and speciesism activation.

In contrast, Wang and Qiu (2024) found that AI technology transparency can have beneficial effects in digital endorser contexts, while Kirkby et al. (2023) found that AI-disclosed brand voice does not necessarily reduce perceived brand authenticity or brand attitude. Liu et al. (2025) further suggest that intelligence disclosure interacts with the temporal phase of the customer journey.

PACT therefore treats disclosure as a four-dimensional architecture encompassing timing, specificity, setting, and product involvement. Timing concerns when consumers are informed about AI involvement. Specificity concerns whether disclosure merely labels the agent as AI or clarifies human oversight, escalation, and data use. Setting distinguishes service interactions from content, endorsement, or brand-voice contexts. Product involvement is treated as a boundary condition because consumers may process AI disclosure differently when a product is personally relevant, risky, or consequential rather than low-involvement or routine.

4.5 Persona selection agency, autonomy, and congruity

Many conversational interfaces use generic agents. A brand featuring multiple AI personas could allow consumers to choose the character they prefer to interact with. According to self-determination theory, autonomy, competence, and relatedness foster engagement and positive experiences (Deci & Ryan, 1985). PACT does not presume that every option is advantageous, because an abundance of choices or poorly differentiated personas can be counterproductive (Iyengar & Lepper, 2000; Scheibehenne et al., 2010). It predicts a bounded agency advantage, in which selecting a familiar persona that is meaningful and manageable makes the interaction feel more self-directed and relationally appropriate. Still, the PACT model must recognize that the desire for a customizable or relational persona is not universal. Some users may rely on a machine heuristic, preferring an AI's perceived objectivity, procedural neutrality, security, or systematic efficiency over anthropomorphic warmth (Sundar & Kim, 2019; Yang & Sundar, 2024). For these consumers, a functional, non-relational interface may foster greater trust than a familiar persona, indicating that persona carryover effects are limited by individual preferences for machine-like neutrality. Self-brand congruity offers an additional pathway through which a chosen persona may be viewed as consistent with a consumer's self-concept, thereby reinforcing identification

and affective trust (Sirgy, 1982; Fournier, 1998). Since autonomy can be immediate while identification often develops through repeated experiences, PACT treats these pathways as related but not interchangeable.

4.6 The cross-touchpoint gap

While the literature offers plausible evidence for the component mechanisms of PACT, it has not yet established the continuity pathway as a whole. Research on virtual influencers typically assesses content exposure, chatbot studies usually examine interactions after they begin, and disclosure research often evaluates cues within a specific context. PACT redirects attention to the transition between settings: whether a consumer perceives the chatbot as the same persona encountered in prior content, transfers parasocial familiarity into the interaction, and treats consistency or inconsistency as consequential for trust. This cross-touchpoint bridge is the model's primary focus and represents a critical empirical challenge for future research.

5. THE PACT MODEL: A STAGE-BY-STAGE CONCEPTUAL ARCHITECTURE

5.1 Organizing logic

The PACT model is organized as a process architecture rather than an assertion of a deterministic consumer response. It proposes that a recognizable AI brand persona can accumulate familiarity during content exposure, carry that familiarity into interaction through perceived continuity, and thereby change the conditions under which trust and brand-related behavior emerge. The model is represented in four stages: Persona Genesis, Familiarity Accumulation, Continuity Transfer, and Trust and Outcomes. The core mechanism is social presence pre-activation: before substantive dialogue begins, recognition of a familiar persona may make the chatbot feel more socially present. However, this is moderated by a temporal trade-off, where functional competence is a prerequisite for initial trust before relational warmth becomes dominant. Pre-activation is a mechanism that gains efficacy as the interaction progresses. The persona carryover effect is expected to operate only above a familiarity threshold and remains vulnerable to persona drift as a continuous technical challenge. Figure 1 presents the parsimonious PACT conceptual research model and maps the eight integrative propositions to the model architecture.

Note

Solid arrows indicate the proposed core paths, dashed arrows indicate qualifying paths, and dotted arrows indicate the negative branch or specificity control. P1-P8 propositions are detailed in Section 6, and all proposed effects remain to be tested.

5.2 Stage 1: Persona Genesis - identity construction

Persona genesis refers to the process by which a brand develops an AI character that maintains consistent recognition over time. Brand value inheritance refers to the perceived alignment between brand identity and persona expression, anthropomorphism calibration involves the extent of human-like visual, linguistic, and behavioral cues, and visual-linguistic identity design provides the recognizable features that facilitate future continuity. At this stage, authenticity stems not only from design but also from continuous management. Practitioners must contend with persona drift, defined here as the technical and experiential risk that an AI agent deviates from its expected identity, tone, values, or conduct across multi-turn conversations. Current technical evidence on persona drift and persona fidelity draws partly on arXiv preprints and should therefore be treated as emerging rather than settled (Brandt, 2025; Brandt & Wang, 2025; Li et al., 2024; Shin et al., 2025).

5.3 Stage 2: Familiarity Accumulation - relational build-up

Through repeated interactions, a consumer may grow accustomed to the persona's tone, narratives, and expected replies. The intensity of parasocial familiarity reflects the strength of this non-reciprocal relational orientation, with perceived persona warmth and narrative immersion proposed as its antecedents. Crucially, this construct differs from brand familiarity, because recognizing a brand's products does not equate to feeling acquainted with its recurring AI character.

5.4 Stage 3: Continuity transfer-recognition in interaction

Continuity transfer begins when consumers encounter the chatbot and recognize it as the same persona previously encountered in feed-based content. Visual, linguistic, narrative, and ethical consistency are expected to facilitate recognition, pre-activate social presence, and make the transition from content to conversation feel coherent. This logic is consistent with expectation-confirmation reasoning: prior expectations become consequential when later experience confirms or disconfirms them (Oliver, 1980). Still, continuity is not guaranteed by reusing the same avatar or brand name. Large language models and role-playing language agents can drift from assigned identity across extended interaction, producing out-of-character behavior, instruction drift, or persona-fidelity failures (Li et al., 2024; Chen et al., 2024; Shin et al., 2025). Persona selection agency may strengthen continuity transfer by making the interaction feel self-selected and relationally appropriate. In contrast, a recognizable but inconsistent persona may violate user expectations more strongly than a neutral, unfamiliar chatbot. Under PACT, this creates a continuity-risk mechanism: the same prior familiarity that enables persona carryover may also intensify negative reactions when the chatbot's tone, competence, values, or disclosure practices conflict with the persona consumers previously encountered.

5.5 Stage 4: Trust and brand outcomes

The interaction leads to assessments of both cognitive and affective trust. Cognitive trust is grounded in perceived competence, accuracy, reliability, transparency, disclosure quality, data security, and appropriate escalation arrangements. Affective trust is grounded in warmth, perceived care, emotional comfort, benevolence, and social presence (Johnson & Grayson, 2005; Glikson & Woolley, 2020; Shang et al., 2024; Wang et al., 2026). PACT does not assume that either trust route always dominates. Instead, it

proposes that prior parasocial familiarity changes the relative salience and sequencing of cognitive and affective trust during the chatbot encounter. For an unfamiliar chatbot, consumers may initially rely more heavily on competence and reliability cues. For a recognized AI brand persona, affective cues such as warmth, familiarity, and relational comfort may become salient earlier, but only if baseline performance is adequate. If the chatbot fails on accuracy, reliability, or escalation, prior familiarity may not protect trust and may even increase disappointment. Downstream outcomes include continued chatbot use, willingness to rely on the chatbot when appropriate, reduced need for human support in low-risk service situations, openness to relevant recommendations, engagement, loyalty, and brand-equity effects. Existing chatbot research supports links between chatbot service quality, cognitive trust, affective trust, satisfaction, and customer loyalty, but PACT extends this logic by asking whether these outcomes are strengthened or weakened when the chatbot is recognized as a previously encountered AI brand persona (Chen et al., 2023; Wang et al., 2026). The core constructs and their proposed observable indicators are summarized in Table 2.

Table 2
Core constructs and operational meaning

Construct	Conceptual meaning	Proposed observable indicators
Brand value inheritance	Perceived embodiment of recognizable brand values, personality, tone, and narrative	Brand-persona coherence; value congruence; authentic value expression; consistency with brand identity
Parasocial familiarity intensity	Relational familiarity built through repeated exposure to the persona in content environments	Felt bond; anticipation of content; perceived familiarity; recognition of persona-specific tone or style
Cross-channel persona recognition	Perception that the chatbot is the same persona previously encountered in content contexts	Same voice; visual identity; narrative continuity; values; interaction style
Social presence pre-activation	Perceived social presence before or at the beginning of the chatbot interaction	Initial warmth; social realness; comfort beginning the conversation; felt presence of a familiar persona
Cognitive trust	Reliability- and competence-based judgement of the chatbot	Accuracy; competence; reliability; transparency; oversight; escalation availability
Affective trust	Relational and emotional confidence in the chatbot or persona	Warmth; benevolence; perceived care; emotional comfort; relational confidence
Continuity violation	Negative disconfirmation that occurs when a recognized persona deviates from expected identity, tone, values, or behavior	Perceived mismatch; out-of-character response; persona drift; disappointment; betrayal-like response; trust decrement
Persona selection agency	Perceived meaningful choice over which AI persona or avatar to interact with	Autonomy; preference-consistent selection; perceived fit; self-congruity; manageable choice set
Disclosure architecture	Configuration of AI disclosure across timing, specificity, setting, and product involvement	Pre-interaction disclosure; clarity of AI identity; human oversight; escalation rights; data-use explanation; high- versus low-involvement product context

6. INTEGRATIVE PROPOSITION DEVELOPMENT

PACT is advanced through eight integrative propositions. The earlier component relationships have not been discarded; rather, they are retained as mediation logic, boundary conditions, controls, or theoretical rationale within the propositions to which they are most closely connected. Prior empirical findings are therefore used as grounding and are not numbered as independent propositions.

6.1 Persona genesis and relational familiarity

P1. Perceived brand value inheritance and calibrated anthropomorphism are expected jointly to strengthen parasocial familiarity with an AI brand persona through perceived authenticity and warmth;

this positive relationship may weaken where humanlike cues become incongruent with the persona's performance capacity or disclosed artificiality.

Theoretical rationale. This proposition preserves the design and authenticity pathways in a single upstream claim. It recognizes that a coherent persona may invite familiarity without assuming that greater realism is always beneficial (Epley et al., 2007; Moore, 2012). It also treats the machine heuristic as a boundary condition: for users who value objectivity, systematic processing, and perceived lack of human bias, calibrated machine-like neutrality may be more trustworthy than relational warmth (Sundar & Kim, 2019; Yang & Sundar, 2024).

6.2 Persona-specific carryover

P2. Relative to equivalent non-persona brand exposure, threshold-level parasocial familiarity with an AI brand persona is expected to increase perceived social presence at the beginning of a later chatbot encounter and reduce perceived interaction friction when the chatbot is recognized as that persona.

Theoretical rationale. P2 retains threshold familiarity and social presence pre-activation. Crucially, competence acts as a functional gateway for the persona carryover effect, as baseline accuracy and reliability are necessary for trust to transition from purely cognitive to affective domains.

6.3 Continuity and its violation

P3. Perceived cross-channel continuity between the content-phase persona and its chatbot instantiation is expected to strengthen the persona carryover effect, whereas recognized discontinuity is expected to generate a continuity violation that may reduce trust below an unfamiliar chatbot baseline.

Theoretical rationale. This proposition consolidates the positive consistency condition and the negative violation prediction. It explicitly includes persona drift as a primary technical driver of continuity violation, where LLM inconsistency triggers disconfirmation of the established relational baseline.

6.4 Dual-route trust conversion and outcomes

P4. In chatbot encounters, social presence pre-activation is expected to increase affective trust, whereas perceived competence and reliability support cognitive trust; further, parasocial familiarity is expected to increase the relative contribution of affective trust to continuance, engagement, and loyalty outcomes.

Theoretical rationale. P4 retains the paper's mediator and trust-route salience claims while avoiding separate propositions for already established trust dimensions. It remains conditional: relational readiness may facilitate trust formation, but does not excuse poor performance (Cheng et al., 2022; Johnson & Grayson, 2005; Prakash & Das, 2024).

6.5 Disclosure architecture

P5. The trust consequences of truthful AI disclosure are expected to depend jointly on timing, specificity, and communication setting: pre-interaction and specific disclosure should better preserve cognitive trust

than delayed or vague disclosure. By contrast, disclosure in moralized or cause-related communication may heighten scepticism, particularly among consumers high in AI aversion.

Theoretical rationale. This integrative proposition preserves the timing, context, specificity, and scepticism logic in one architecture rather than treating each design dimension as an independent theoretical contribution (Luo et al., 2019; Wang & Qiu, 2024; De Cicco et al., 2025; Liu et al., 2025; Qiu et al., 2025).

6.6 Performance-expectation fit

P6. The positive influence of humanlike persona cues on cognitive trust is expected to depend on chatbot competence and reliability, such that salient service failure produces a sharper trust decrement where a highly anthropomorphic or already familiar persona has raised consumer expectations.

Theoretical rationale. P6 preserves the competence mediation and mismatch penalty while expressing them as one performance-contingent proposition. It guards against interpreting warmth or familiarity as a substitute for functional quality.

6.7 Consumer persona selection agency

P7. Meaningful selection of a familiar AI persona, rather than assignment to an otherwise equivalent avatar, is expected to increase affective trust, engagement and continuance through perceived autonomy and self-brand congruity.

Theoretical rationale. This proposition retains both agency and congruity mechanisms within a manageable claim. It treats choice as bounded: agency is likely to be beneficial only where options are intelligible, non-manipulative, and relevant to the consumer (Deci & Ryan, 1985; Fournier, 1998; Sirgy, 1982).

6.8 Digital-literacy boundary condition

P8. Consumer digital literacy is expected to moderate PACT trust formation, strengthening the influence of specific transparency and competence information on cognitive trust while attenuating reliance on anthropomorphic social cues alone.

Theoretical rationale. P8 retains a theoretically meaningful consumer-level moderator without multiplying propositions for every contextual variation. Cross-cultural, linguistic, and service-risk variations remain appropriate boundary tests in the empirical programme rather than additional propositions (van Deursen & van Dijk, 2011; Petty & Cacioppo, 1986).

7. PROPOSED EMPIRICAL VALIDATION PROGRAM

7.1 Overall research design

PACT requires cumulative validation. To align with GenAI-era research standards, experimental validation should include multi-turn consistency assessment for persona adherence, with LLM-as-a-judge methods used only as auxiliary measures after validation against human ratings.

The critical test is not simply whether a familiar avatar produces more positive ratings. It is whether persona-specific familiarity produces a continuity advantage beyond equivalent brand familiarity, whether social presence mediates that advantage, and whether an inconsistent persona reverses it. Study 3 is consequently the pivotal design in the program.

The proposed empirical validation roadmap is outlined in Table 3.

Table 3
Four-study validation roadmap

Study	Design and stimulus	Principal test	Outcome domains	Analytic emphasis
1. Persona cues and trust	Experiment manipulating brand-value inheritance and anthropomorphism calibration, with product involvement varied through pretested high-versus low-involvement product categories	Whether coherent persona cues and calibrated anthropomorphism increase authenticity, warmth, parasocial familiarity, and trust	Authenticity, warmth, parasocial familiarity, social presence, cognitive trust, affective trust, purchase or continuance intention	Non-linear contrasts; mediation; moderated mediation by product involvement
2. Disclosure architecture	Experiment manipulating AI disclosure timing, disclosure specificity, communication setting, and product involvement; AI aversion measured as a consumer-level moderator	Whether disclosure effects differ by timing, specificity, setting, product involvement, and AI aversion	Cognitive trust, affective trust, skepticism, persuasion knowledge, firm trust, ad attitude, continuance intention	Conditional indirect effects; moderated mediation; planned contrasts between specific versus vague disclosure
3. Continuity transfer	Two-stage exposure experiment with delay: persona-based content versus non-persona brand content, followed by consistent versus inconsistent chatbot interaction, with choice versus assignment as an additional factor	Whether prior persona familiarity transfers into chatbot interaction, and whether inconsistency creates continuity violation	Cross-channel recognition, pre-chat social presence, interaction friction, cognitive trust, affective trust, switching intention, recommendation receptivity	Planned contrasts; sequential mediation; human-coded persona consistency; auxiliary automated consistency metrics; LLM-as-a-judge only after validation against human ratings
4. Field and cross-cultural validation	Survey and behavioral follow-up among real followers in India and an East Asian comparison market, such as China, using	Nomological validity, contextual robustness, and cross-cultural boundary conditions	Actual continuance, escalation behavior, engagement, loyalty, recommendation	Measurement invariance; SEM; multi-group comparison; out-of-sample prediction;

Note. Study 3 should not rely solely on LLM-as-a-judge scoring. Automated consistency scores should be validated against human-coded persona consistency and participant recognition checks. "Prompt-to-Line Consistency" should be defined as a study-specific metric unless an established source is provided. Study 4 requires comparable samples, localized measures, and measurement-invariance testing before cross-market differences are interpreted.

Study	Design and stimulus	Principal test	Outcome domains	Analytic emphasis
	comparable brand/persona contexts and localized instruments		acceptance, brand equity	behavioral validation
<i>Note. Study 3 should not rely solely on LLM-as-a-judge scoring. Automated consistency scores should be validated against human-coded persona consistency and participant recognition checks. “Prompt-to-Line Consistency” should be defined as a study-specific metric unless an established source is provided. Study 4 requires comparable samples, localized measures, and measurement-invariance testing before cross-market differences are interpreted.</i>				

7.2 The decisive sequential continuity experiment

Study 3 should expose participants to a branded content stream through repeated encounters instead of a single vignette. Participants would be randomly assigned to either an AI persona content condition or an informational brand-content control condition, with the latter matched on brand, product, message volume, and visual quality. After a sufficient interval, they should engage in a service task via either a persona-consistent chatbot or one that is deliberately inconsistent. A distinct choice manipulation can compare the selection of a persona with assignment to the same persona. Before interpreting outcome tests, manipulation checks must evaluate the intensity of parasocial familiarity, brand familiarity, cross-channel recognition, perceived persona consistency, and the clarity of perceived AI disclosure. Persona carryover will be supported only if the avatar exposure condition demonstrates significantly higher initial social presence and lower interaction friction than the non-avatar brand control, thereby exerting an indirect influence on affective trust and behavioral outcomes. A continuity violation would require a planned contrast in which the familiar yet inconsistent condition is trusted less than an unfamiliar chatbot baseline. Since brief studies may fail to generate sufficient parasocial familiarity (Boerman & van Reijmersdal, 2020), a non-significant result should not be automatically interpreted as disconfirmation. A null result stemming from weak familiarity manipulation would mainly indicate a design flaw, whereas a null result derived from validated high familiarity, recognition, and sufficient statistical power would offer greater theoretical insight.

7.3 Measurement and analysis considerations

Wherever feasible, established measures should be employed to assess cognitive and affective trust, parasocial interaction, social presence, engagement, and loyalty. To enhance empirical rigor and reduce potential bias in self-reported data, studies may supplement traditional psychometric scales with automated, objective multi-turn evaluation frameworks, including validated LLM-as-a-judge scoring, to assess Prompt-to-Line Consistency and verify persona adherence throughout the dialogue. New or significantly adapted scales are necessary to measure brand value inheritance, cross-channel persona recognition, perceived interaction friction, disclosure specificity, and persona selection agency. These measures should initially undergo expert review, cognitive interviews, and a pilot confirmatory factor

analysis. When constructs shift from human media figures to explicitly disclosed AI personae, the wording must avoid implying that consumers believe the agent is sentient or genuinely caring.

Additionally, multi-group analysis (MGA) within the structural equation modeling (SEM) framework can test statistically significant differences across consumer segments. This approach is appropriate for validating moderating roles such as digital literacy and cultural background by comparing path coefficients across groups, provided measurement invariance is established first.

The empirical program should incorporate longitudinal design and multi-turn evaluation to test whether persona continuity and trust formation remain stable across repeated interactions. Automated evaluation, including LLM-based scoring, may be used as an auxiliary tool for assessing persona consistency, but such scores should be validated against human-coded consistency ratings and participant recognition checks.

The statistical method should be selected after confirming measurement specifications. Covariance-based SEM is appropriate when the study is confirmatory and constructs are modeled reflectively with adequate distributional properties. PLS-SEM may be appropriate when prediction-oriented objectives, complex models, or well-justified composite measures dominate (Hair et al., 2021). In either case, studies should report measurement validity, indirect-effect confidence intervals, robustness checks, model comparisons, and preregistered contrasts. Multi-group comparisons should be preceded by measurement-invariance tests.

To support credible inference, experiments should be preregistered, include power analyses suited to moderated mediation and planned contrasts, and report exclusions, attrition, and unsuccessful manipulations. Field studies should collect behavioral evidence, such as chatbot continuation or human escalation, only with informed consent and proportionate data protection.

7.4 Indian and cross-cultural validation opportunity

The Indian digital consumer setting is theoretically valuable because brands increasingly operate in a large, mobile-first, multilingual, and culturally heterogeneous interaction environment. Recent evidence from the Internet and Mobile Association of India and Kantar (2025) indicates that India has one of the world's largest active internet populations and that Indic-language use is central to digital consumption. India should nevertheless be treated as a validation context rather than presumed evidence for PACT. A robust design would recruit participants in India and at least one comparison market, test measurement equivalence, examine language and platform familiarity, and assess whether trust pathways, disclosure expectations, or persona-selection benefits differ across settings.

8. DISCUSSION

8.1 Theoretical contribution

PACT's main contribution lies in its proposed shift in the unit of analysis. While existing research offers valuable insights into virtual influencers, chatbot trust, and disclosure, consumers are increasingly encountering these elements not as isolated applications but as a unified, recognizable AI brand persona throughout their journey. By theorizing the shift from feed to conversation, PACT introduces temporal and cross-channel structure to the formation of AI trust. The model's adaptability to diverse cultural contexts also presents an important area for further investigation. In markets such as East Asia, trust building may integrate symbolic elements, including cultural alignment and collective identity, alongside the affective and cognitive pathways outlined in the PACT model. Identifying these symbolic elements is a potential expansion of the framework, thereby strengthening its applicability to global digital branding strategies.

Second, the persona-carryover and continuity-violation propositions render continuity empirically significant. Familiarity is not always advantageous. Maintaining a coherent and capable persona can reduce interaction friction and accelerate affective trust, whereas inconsistency or inadequate performance may lead to heightened disappointment. This symmetry enhances the model's falsifiability and ethical significance. The technical challenge is that large language models can unintentionally deviate from assigned personas during multi-turn conversations. Such drift can prevent the model from reaching the intended continuity stage because even minor inconsistencies may provoke the continuity violation effect and erode relational equity.

Third, PACT incorporates disclosure findings while avoiding the reduction of transparency to a mere cost of trust. PACT suggests maintaining truthful disclosure while testing timing, specificity, and communication setting as design features. This approach aligns with consumer autonomy, as trust built on continuity should not rely on concealing artificial identities or gathering relational data without explanation.

Fourth, the consumer persona selection agency expands AI brand relationship thinking beyond the passive assignment of a bot. By connecting meaningful choice to autonomy and self-brand congruity, PACT initiates a research agenda exploring the conditions under which consumer-controlled relational interfaces can foster trust without devolving into manipulative personalization. However, a balanced perspective must acknowledge the machine heuristic. Some users may prefer non-personified AI because they perceive it as more objective, systematic, and less susceptible to human-like bias (Sundar & Kim, 2019; Yang & Sundar, 2024). This argues that extensive personalization is not always beneficial, offering a more refined perspective for future studies on user preferences between machine-like neutrality and anthropomorphic warmth.

8.2 Managerial and governance implications

Managers may be tempted to interpret a popular AI persona as an asset that can be deployed indiscriminately across all consumer touchpoints. PACT suggests greater caution. Where the same persona moves from content to service, design governance should preserve recognizable voice and values while acknowledging the technical challenge of persona drift. Practitioners must actively manage

this drift to keep an AI consistently in character and ensure that promised warmth is supported by competent service capabilities and accessible human escalation.

Disclosure

should occur before substantive interaction and should be understandable rather than merely legalistic. However, practitioners must consider the potential differential impacts of disclosure based on product involvement. For high-involvement products, explicit AI labeling can lead to a sharper decline in trust compared to low-involvement categories. Consumers should be informed that they are interacting with an AI persona, whether a human can intervene, how data generated in the conversation may be used, and how to opt for an alternative channel. For a familiar persona, disclosure should also avoid implying a human memory or relationship that the system does not possess.

Choice architecture requires similar restraint. Offering a small number of clearly differentiated personas may support autonomy; nudging consumers toward a persona based on opaque profiling could undermine it. Brands should therefore treat persona-selection data as sensitive interaction preference information and avoid inferring vulnerability, identity or emotion without appropriate permission and safeguards.

8.3 Limitations and directions for future research

The proposed model has several limitations. First, it is conceptual and therefore does not demonstrate that persona carryover exists, reaches a practically important magnitude, or generalizes across products and cultures. Second, its content-first pathway may not describe journeys in which consumers initially meet a chatbot and later follow the persona in content environments. Reverse or reciprocal transfer should be examined rather than assumed.

Third, parasocial familiarity is difficult to induce ethically and ecologically in short experiments. Longitudinal content exposure and studies of real followers may be required. Fourth, high-stakes service contexts such as health, finance, or complaints make product involvement a critical moderator; for these high-involvement products, users may exhibit a strong human-labeling bias that limits or negates the affective advantage of familiarity. Fifth, personas may generate adverse outcomes, including dependency, stereotyping, identity manipulation, privacy intrusion, or reluctance to seek human assistance.

Future research should examine trust repair after persona failure, the effects of language and cultural relational norms, the point at which persona choice becomes burdensome, consumer responses to continuity disclosure, and the behavior of vulnerable or younger users. Importantly, studies should evaluate not only engagement and purchase intentions but also technical solutions for mitigating persona drift and maintaining identity consistency, alongside informed choice, appropriate escalation, and consumer welfare.

9. CONCLUSION

AI brand personas may increasingly accompany consumers from content consumption into conversation. PACT proposes that this transition matters: parasocial familiarity developed in a feed may pre-activate social presence when the same persona appears as a chatbot, thereby influencing how trust begins to form. The proposed effect is neither automatic nor inherently beneficial; it is conditional on recognizable continuity, competence, transparent disclosure, and meaningful consumer control.

Taken together, the model answers RQ1 by proposing that transfer is most likely when persona familiarity exceeds a meaningful threshold and the chatbot is recognized as the same persona. It answers RQ2 by specifying social presence pre-activation as the bridge from familiarity to affective trust, while competence and reliability sustain cognitive trust. It answers RQ3 by identifying continuity, disclosure architecture, performance-expectation fit, persona selection agency, machine-heuristic preferences, digital literacy, product involvement, and culture as conditions that can strengthen, weaken, or reverse the proposed effects.

By grounding eight integrative propositions in relevant prior evidence and setting out a staged validation program, this paper offers a cautious but testable framework for studying continuous AI brand relationships. Its central normative position is equally important: if brands design continuity as a trust resource, they should govern that continuity transparently and preserve the consumer's ability to understand, choose, and escalate the interaction.

Declarations

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of interest

The authors declare no known competing financial interests or personal relationships that could have influenced the work reported in this article.

Data availability

Data sharing is not applicable to this article because no new empirical data were generated or analyzed.

Ethics approval

This article is a conceptual theory-building paper and does not involve human participants, human data, or animal subjects. Therefore, ethics approval was not required.

Use of generative AI and AI-assisted technologies

Generative AI tools, including Claude by Anthropic and ChatGPT by OpenAI, were used only for language refinement, editing support, and clarity checking. The theoretical development, interpretation, argumentation, and intellectual content of the manuscript are the authors' own work. The AI tools were not used to generate empirical findings, conduct statistical analysis, or make independent scholarly interpretations.

Author Contribution

Rohan Kumar HM conceived the study, conducted the literature review and conceptual synthesis, developed the Persona-Avatar Continuity-Trust (PACT) Model and its propositions, prepared the figure and tables, and wrote the original manuscript. Dr. Clement Sudhahar contributed to the theoretical framing, refinement of the conceptual model and propositions, development of the proposed empirical validation programme, and critical revision of the manuscript. Dr. Sumanjit Das supervised the study, contributed to the conceptual development and interpretation of the framework, reviewed and edited the manuscript, and coordinated the submission process. All authors reviewed and approved the final manuscript.

References

1. Allal-Chérif O, Puertas R, Carracedo P (2024) Intelligent influencer marketing: How AI-powered virtual influencers outperform human influencers. *Technol Forecast Soc Chang* 200 Article 123113. <https://doi.org/10.1016/j.techfore.2023.123113>
2. Boerman SC, van Reijmersdal EA (2020) Disclosing influencer marketing on YouTube to children: The moderating role of parasocial relationship. *Front Psychol* 10., Article 3042. <https://doi.org/10.3389/fpsyg.2019.03042>
3. Brandt TJ (2025) *Navigating the synchrony-stability frontier in adaptive chatbots*. arXiv. <https://doi.org/10.48550/arXiv.2510.00339>
4. Brandt TJ, Wang CX (2025) *The adaptation paradox: Agency vs. mimicry in companion chatbots*. arXiv. <https://doi.org/10.48550/arXiv.2509.12525>
5. Chen J, Wang X, Xu R, Yuan S, Zhang Y, Shi W, Xie J, Li S, Yang R, Zhu T, Chen A, Li N, Chen L, Hu C, Wu S, Ren S, Fu Z, Xiao Y (2024) *From persona to personalization: A survey on role-playing language agents*. arXiv. <https://doi.org/10.48550/arXiv.2404.18231>
6. Chen Q, Lu Y, Gong Y, Xiong J (2023) Can AI chatbots help retain customers? Impact of AI service quality on customer loyalty. *Internet Res* 33(6):2205–2243. <https://doi.org/10.1108/INTR-09-2021-0686>
7. Cheng X, Bao Y, Zarifis A, Gong W, Mou J (2022) Exploring consumers' response to text-based chatbots in e-commerce: The moderating role of task complexity and chatbot disclosure. *Internet Res* 32(2):496–517. <https://doi.org/10.1108/INTR-08-2020-0460>

8. Dabiran E, Farivar S, Wang F, Grant G (2024) Virtually human: Anthropomorphism in virtual influencer marketing. *J Retailing Consumer Serv* 79 Article 103797. <https://doi.org/10.1016/j.jretconser.2024.103797>
9. De Cicco R, Elmashhara MG, Silva SC, Hammerschmidt M (2025) The impact of providing non-human identity cues about sales agents on consumer responses: The role of social presence and speciesism activation. *Eur J Mark* 59(13):55–84. <https://doi.org/10.1108/EJM-01-2022-0066>
10. De Cicco R, Iacobucci S, Cannito L, Onesti G, Ceccato I, Palumbo R (2024) Virtual vs. human influencer: Effects on users' perceptions and brand outcomes. *Technol Soc* 77:102488. <https://doi.org/10.1016/j.techsoc.2024.102488>
11. De Cicco R, Silva SC, Palumbo R (2021) Should a chatbot disclose itself? Implications for an online conversational retailer. In A. Følstad, T. Araujo, S. Papadopoulos, E. L.-C. Law, O.-C. Granmo, E. Luger, & P. B. Brandtzaeg (Eds.), *Chatbot research and design* (pp. 3–15). Springer. https://doi.org/10.1007/978-3-030-68288-0_1
12. Deci EL, Ryan RM (1985) *Intrinsic motivation and self-determination in human behavior*. Plenum. <https://doi.org/10.1007/978-1-4899-2271-7>
13. Dunn JR, Schweitzer ME (2005) Feeling and believing: The influence of emotion on trust. *J Personal Soc Psychol* 88(5):736–748. <https://doi.org/10.1037/0022-3514.88.5.736>
14. Epley N, Waytz A, Cacioppo JT (2007) On seeing human: A three-factor theory of anthropomorphism. *Psychol Rev* 114(4):864–886. <https://doi.org/10.1037/0033-295X.114.4.864>
15. Fournier S (1998) Consumers and their brands: Developing relationship theory in consumer research. *J Consum Res* 24(4):343–373. <https://doi.org/10.1086/209515>
16. Gefen D, Straub DW (2004) Consumer trust in B2C e-commerce and the importance of social presence: Experiments in e-products and e-services. *Omega* 32(6):407–424. <https://doi.org/10.1016/j.omega.2004.01.006>
17. Gilson LL, Goldberg CB (eds) (2015) Editors' comment: So, what is a conceptual paper? *Group & Organization Management*, 40(2), 127–130. <https://doi.org/10.1177/1059601115576425>
18. Glikson E, Woolley AW (2020) Human trust in artificial intelligence: Review of empirical research. *Acad Manag Ann* 14(2):627–660. <https://doi.org/10.5465/annals.2018.0057>
19. Goel P, Garg A (2025) Virtual personalities, real bonds: Anthropomorphised virtual influencers' impact on trust and engagement. *J Consumer Mark* 42(4):431–447. <https://doi.org/10.1108/JCM-05-2024-6915>
20. Gutuleac R, Baima G, Rizzo C, Bresciani S (2024) Will virtual influencers overcome the uncanny valley? The moderating role of social cues. *Psychol Mark* 41(7):1419–1431. <https://doi.org/10.1002/mar.21989>
21. Hair JF Jr., Hult GTM, Ringle CM, Sarstedt M, Danks NP, Ray S (2021) *Partial least squares structural equation modeling (PLS-SEM) using R: A workbook*. Springer. <https://doi.org/10.1007/978-3-030-80519-7>

22. Hoffner CA, Bond BJ (2022) Parasocial relationships, social media, & well-being. *Curr Opin Psychol* 45:101306. <https://doi.org/10.1016/j.copsyc.2022.101306>
23. Horton D, Wohl RR (1956) Mass communication and parasocial interaction: Observations on intimacy at a distance. *Psychiatry* 19(3):215–229. <https://doi.org/10.1080/00332747.1956.11023049>
24. Huang R, Kim M, Lennon S (2022) Trust as a second-order construct: Investigating the relationship between consumers and virtual agents. *Telematics Inform* 70 Article 101811. <https://doi.org/10.1016/j.tele.2022.101811>
25. Internet and Mobile Association of India, & Kantar (2025) Internet in India 2024: Kantar–IAMAI report. Internet and Mobile Association of India
26. Iyengar SS, Lepper MR (2000) When choice is demotivating: Can one desire too much of a good thing? *J Personal Soc Psychol* 79(6):995–1006. <https://doi.org/10.1037/0022-3514.79.6.995>
27. Jaakkola E (2020) Designing conceptual articles: Four approaches. *AMS Rev* 10(1–2):18–26. <https://doi.org/10.1007/s13162-020-00161-0>
28. Johnson D, Grayson K (2005) Cognitive and affective trust in service relationships. *J Bus Res* 58(4):500–507. [https://doi.org/10.1016/S0148-2963\(03\)00140-1](https://doi.org/10.1016/S0148-2963(03)00140-1)
29. Kirkby A, Baumgarth C, Henseler J (2023) To disclose or not disclose, is no longer the question: Effect of AI-disclosed brand voice on brand authenticity and attitude. *J Prod Brand Manage* 32(7):1108–1122. <https://doi.org/10.1108/JPBM-02-2022-3864>
30. Kull AJ, Romero M, Monahan L (2021) How may I help you? Driving brand engagement through the warmth of an initial chatbot message. *J Bus Res* 135:840–850. <https://doi.org/10.1016/j.jbusres.2021.03.005>
31. Li K, Liu T, Bashkansky N, Bau D, Viégas F, Pfister H, Wattenberg M (2024) *Measuring and controlling persona drift in language model dialogs*. arXiv. <https://doi.org/10.48550/arXiv.2402.10962>
32. Liu N, Shi S, Huang Y, Zhang M, Leung WKS (2025) From attraction to retention: The influence of AI chatbot intelligence disclosure on customer interaction intention. *Int J Contemp Hospitality Manage* 37(8):2581–2600. <https://doi.org/10.1108/IJCHM-05-2024-0766>
33. Lou C (2022) Social media influencers and followers: Theorization of a trans-parasocial relation and explication of its implications for influencer advertising. *J Advertising* 51(1):4–21. <https://doi.org/10.1080/00913367.2021.1880345>
34. Lou C, Kiew STJ, Chen T, Lee TYM, Ong JEC, Phua Z (2023) Authentically fake? How consumers respond to the influence of virtual influencers. *J Advertising* 52(4):540–557. <https://doi.org/10.1080/00913367.2022.2149641>
35. Luo X, Tong S, Fang Z, Qu Z (2019) Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Mark Sci* 38(6):937–947. <https://doi.org/10.1287/mksc.2019.1192>
36. MacInnis DJ (2011) A framework for conceptual contributions in marketing. *J Mark* 75(4):136–154. <https://doi.org/10.1509/jmkg.75.4.136>

37. McAllister DJ (1995) Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations. *Acad Manag J* 38(1):24–59. <https://doi.org/10.5465/256727>
38. McKnight DH, Chervany NL (2001) What trust means in e-commerce customer relationships: An interdisciplinary conceptual typology. *Int J Electron Commer* 6(2):35–59. <https://doi.org/10.1080/10864415.2001.11044235>
39. Moore RK (2012) A Bayesian explanation of the uncanny valley effect and related psychological phenomena. *Sci Rep* 2:864. <https://doi.org/10.1038/srep00864>
40. Mori M, MacDorman KF, Kageki N (2012) The uncanny valley. *IEEE Rob Autom Magazine* 19(2):98–100. <https://doi.org/10.1109/MRA.2012.2192811>
41. Mozafari N, Weiger WH, Hammerschmidt M (2022) Trust me, I'm a bot: Repercussions of chatbot disclosure in different service frontline settings. *J Service Manage* 33(2):221–245. <https://doi.org/10.1108/JOSM-10-2020-0380>
42. Nordheim CB, Følstad A, Bjørkli CA (2019) An initial model of trust in chatbots for customer service: Findings from a questionnaire study. *Interact Comput* 31(3):317–335. <https://doi.org/10.1093/iwc/iwz022>
43. Oliver RL (1980) A cognitive model of the antecedents and consequences of satisfaction decisions. *J Mark Res* 17(4):460–469. <https://doi.org/10.1177/002224378001700405>
44. Petty RE, Cacioppo JT (1986) Communication and persuasion: Central and peripheral routes to attitude change. Springer
45. Prakash AV, Das S (2024) (Why) do we trust AI? A case of AI-based health chatbots. *Australasian J Inform Syst.* 28<https://doi.org/10.3127/ajis.v28.4235>
46. Qiu X, Wang Y, Zeng Y, Cong R (2025) Artificial intelligence disclosure in cause-related marketing: A persuasion knowledge perspective. *J Theoretical Appl Electron Commer Res* 20(3):193. <https://doi.org/10.3390/jtaer20030193>
47. Rubin AM, Perse EM, Powell RA (1985) Loneliness, parasocial interaction, and local television news viewing. *Hum Commun Res* 12(2):155–180. <https://doi.org/10.1111/j.1468-2958.1985.tb00071.x>
48. Scheibehenne B, Greifeneder R, Todd PM (2010) Can there ever be too many options? A meta-analytic review of choice overload. *J Consum Res* 37(3):409–425. <https://doi.org/10.1086/651235>
49. Shang R, Hsieh G, Shah C (2024) Trusting your AI agent emotionally and cognitively: Development and validation of a semantic differential scale for AI trust. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 7, 1343–1356
50. Shin J, Oh J, Kim E, Song H, Oh A (2025) *Spotting out-of-character behavior: Atomic-level evaluation of persona fidelity in open-ended generation*. arXiv. <https://doi.org/10.48550/arXiv.2506.19352>
51. Sirgy MJ (1982) Self-concept in consumer behavior: A critical review. *J Consum Res* 9(3):287–300. <https://doi.org/10.1086/208924>
52. Song L, Chow OW, Na M, Sarjit SG (2024) Virtual idol and youth identity: The impact of mainstream culture on the virtual idol online fanbase community in China. *Stud Media Communication*

12(2):100–115. <https://doi.org/10.11114/smc.v12i2.6658>

53. Stein J-P, Breves PL, Anders N (2024) Parasocial interactions with real and virtual influencers: The role of perceived similarity and human-likeness. *New Media Soc* 26(6):3433–3453. <https://doi.org/10.1177/14614448221102900>
54. Sundar SS, Kim J (2019) Machine heuristic: When we trust computers more than humans with our personal information. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery. <https://doi.org/10.1145/3290605.3300768>
55. Uysal E, Alavi S, Bezençon V (2022) Trojan horse or useful helper? A relationship perspective on artificial intelligence assistants with humanlike features. *J Acad Mark Sci* 50:1153–1175. <https://doi.org/10.1007/s11747-022-00856-9>
56. van Deursen AJAM, van Dijk JAGM (2011) Internet skills and the digital divide. *New Media Soc* 13(6):893–911. <https://doi.org/10.1177/1461444810386774>
57. Wang S, Fatima N, Shahbaz M, Asif M (2026) Building user trust in AI chatbots for customer service through human-like cues and perceived reliability. *Sci Rep* 16:7860. <https://doi.org/10.1038/s41598-026-38179-2>
58. Wang X, Qiu X (2024) The positive effect of artificial intelligence technology transparency on digital endorsers: Based on the theory of mind perception. *J Retailing Consumer Serv* 78:103777. <https://doi.org/10.1016/j.jretconser.2024.103777>
59. Yang H, Sundar SS (2024) Machine heuristic: Concept explication and development of a measurement scale. *J Computer-Mediated Communication* 29(6) Article zmae019. <https://doi.org/10.1093/jcmc/zmae019>

Figures




Figure 1

Parsimonious Persona-Avatar Continuity-Trust (PACT) conceptual research model: eight integrative propositions.

Note: Solid arrows indicate the proposed core paths, dashed arrows indicate qualifying paths, and dotted arrows indicate the negative branch or specificity control. P1-P8 propositions are detailed in Section 6, and all proposed effects remain to be tested.