Information & Management 64 (2027) 104465

Contents lists available at ScienceDirect

Information & Management

journal homepage: www.elsevier.com/locate/im

AI error or AI creativity? Attributional ambiguity in AI-generated expectation-violating prerolls

Shichen Zhang [a,b], Liang Peng [c], Yinliang (Ricky) Tan [d,*]

[a] College of Management and Economics, Tianjin University, Tianjin, China
[b] Laboratory of Computation and Analytics of Complex Management Systems, Tianjin University, Tianjin, China
[c] Youju Interactive Beijing Technology Co. Ltd, Beijing, China
[d] China Europe International Business School, Shanghai, China

ARTICLE INFO

Keywords:
Generative AI
Attributional ambiguity
Preroll advertising
Cognitive processing
Field experiment

ABSTRACT

Generative AI has sharply reduced the cost of producing and experimenting with short-video advertising creatives, making AI-mediated content production an increasingly central component of digital platform ecosystems. Yet how different forms of AI-generated content shape user cognitive processing and downstream behavioral responses remains poorly understood. AI-generated violations differ from traditional human-crafted incongruity because they introduce *attributional ambiguity*: Viewers must infer whether an observed oddity reflects intentional AI creativity or generative-model error. Drawing on dual-stage processing as a cognitive foundation, we propose that the resolution of attributional ambiguity depends on the processing demands foregrounded by the violation. Image violations foreground the perceptual coherence of an object's appearance, potentially favoring an AI-error interpretation that reduces fluency and credibility. Behavior violations foreground the interpretation of an actor–action relationship, potentially favoring intentional AI creativity and sustaining curiosity and attention. Using a field experiment, we show that behavior-violating prerolls increase sustained engagement, whereas the corresponding effect of image-violating prerolls is not statistically significant. Both forms of AI-generated violations positively affect conversion behavior. AI disclosure produces opposing interaction effects on sustained engagement for image and behavior violations, consistent with easing interpretation of unusual imagery but dampening narrative curiosity, respectively. Hedonic products weaken the effects of both violation types, consistent with hedonic crowding-out. This study identifies attributional ambiguity as a distinctive mechanism underlying responses to AI-generated content and advances understanding of how generative AI reshapes advertising effectiveness on digital platforms.

1. Introduction

The development of the digital economy has intensified the imbalance between users' limited cognitive resources and the rapidly growing volume of online information. As information overload becomes a defining feature of digital economy [1,2], attention increasingly functions as the scarce resource in marketing [3,4]. Short-video platforms, characterized by the fast pace, fragmented formats, and algorithmic content distribution, further amplify this scarcity while simultaneously creating new opportunities for brand communication [5,6]. Yet traditional advertising, which relies heavily on manual content creation, struggles to maintain effectiveness amid the rapidly evolving digital economy.

Recent industry analyses highlight generative AI (GenAI) as a major driver of change in creative production and digital marketing. GenAI can boost creative productivity, support real-time optimization, and enhance user interaction across the customer journey [7,8]. Companies such as Coca-Cola and Unilever are already using GenAI for product cocreation and automated ad development, as illustrated by Coca-Cola's AI-enabled Y3000 beverage line. Using large language models, diffusion models, and multimodal generation algorithms [9,10], AI systems can create unexpected visuals or actions (e.g., levitating objects, physics-defying motion) that immediately draw user attention. These developments reflect a broader shift from human-driven "creative logic" to algorithm-driven "AI logic", which primarily unfolds in short-video platforms [11].

A notable application of AI logic is the rise of *attention-grabbing prerolls*, defined as very short, often product-unrelated clips inserted immediately before the main advertisement (see Fig. 1). These prerolls typically introduce brief visual or behavioral surprises that interrupt

* Corresponding author.
E-mail addresses: sczhang@tju.edu.cn (S. Zhang), ccid@263.net (L. Peng), yrtan@ceibs.edu (Y. Tan).

https://doi.org/10.1016/j.im.2026.104465
Received 1 February 2026; Received in revised form 31 October 2026; Accepted 2 November 2026
Available online 9 November 2026
0378-7206/© 2026 Published by Elsevier B.V.

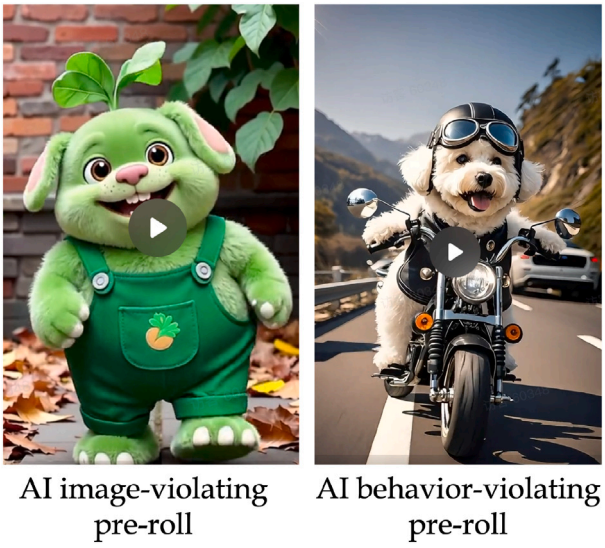


Fig. 1. AI preroll examples.
Notes. Image-violating prerolls depict incongruity in an object's visual appearance or material composition (e.g., a dog composed of vegetables). Behavior-violating prerolls depict incongruity in the relationship between an actor and its action (e.g., a dog riding a motorcycle).

users' habitual scrolling patterns and prime attention for subsequent ad exposure. In this sense, they function as an attentional gating mechanism that disrupts passive browsing and increases the likelihood of full exposure to the commercial message [12]. Major short-video platforms have begun large-scale deployment of generative-AI advertising infrastructures (e.g., TikTok Symphony, Meta Advantage+, and Google AI ad tools), reflecting a broader shift toward AI-generated advertising content [13–15]. Enabled by these generative-AI systems, advertisers can rapidly generate and iteratively test large volumes of such clips, supporting continuous experimentation and performance optimization [16]. In our empirical setting, the collaborating firm — a leading digital advertising agency holding approximately 16.0% market share in China's short-video e-commerce advertising sector — has also widely adopted this format in real campaign deployments.

Despite the practical application of AI-generated prerolls across short-video platforms, how this new mode of AI-mediated content production reshapes users' cognitive processing and behavioral responses on digital platforms remains poorly understood. We distinguish two forms of AI-generated prerolls according to the structure of the expectation violation. *Image-violating* prerolls violate expectations about an object's visual appearance or material composition, such as a dog composed of vegetables; *behavior-violating* prerolls violate expectations about an actor's actions or capabilities, such as a dog riding a motorcycle (Fig. 1). That is, image violations concern the attributes constituting an object's visual representation, whereas behavior violations concern the relationship between an actor and an action within an event. Whether, why, and under what conditions these two forms elicit different user responses is an open empirical question that the information systems (IS) literature has not yet systematically addressed; given the rapid platform-side rollout of generative AI advertising suites by short-video platforms [13–15], addressing it is increasingly pressing.

A natural starting point is the long tradition of work on incongruity and surprise in advertising [17–22], which explains why expectation violations capture attention. This literature, however, implicitly assumes that violations originate from intentional human creative design. AI-generated content disrupts this assumption: When a viewer encounters an oddity in an AI-produced preroll, she must evaluate not only what the violation means but also why it is there—*whether it reflects deliberate creative intent or an unintended generative-model error*. We refer to uncertainty between these competing interpretations as *attributional ambiguity*. Connecting this construct to the algorithm-aversion and algorithm-appreciation literature [23–27], we propose attributional ambiguity as the central mechanism that differentiates AI-generated from human-crafted incongruity and that may underlie heterogeneous user responses to AI-generated content. We next consider how the structure of the violation shapes the resolution of this ambiguity.

If attributional ambiguity is the source of heterogeneity, the direction in which it is resolved should depend on the cognitive processing that the violation foregrounds. Drawing on dual-stage processing as a cognitive foundation [28–31], we propose that image-violating prerolls foreground an object's visual configuration, with unusual visual features providing early cues to incongruity during preattentive processing. Interpreting this incongruity centers on reconciling the object's appearance with its familiar representation. This emphasis on perceptual coherence makes the quality of the generated visual representation salient, which may favor an attribution to generative-model error, potentially eroding processing fluency and credibility. Behavior-violating prerolls, in contrast, foreground an actor–action relationship: Understanding the violation requires relating what the actor does to expectations about its capabilities and the depicted situation. This relational interpretation engages narrative processing and provides a basis for causal questions (why is the dog riding a motorcycle?), potentially favoring an attribution to intentional AI creativity that sustains curiosity and prolongs engagement. Thus the two forms foreground different bases for resolving attributional ambiguity: the perceptual coherence of an object's appearance for image violations and the interpretation of an actor–action relationship for behavior violations. How these different processing demands shape attributional interpretation and subsequent user responses has not been systematically examined in the IS literature on AI-mediated content. This relationship motivates the first empirical question this study addresses.

Attributional ambiguity also has direct implications for two contextual factors that shape how viewers respond to AI content on digital platforms. The first is AI disclosure: An explicit AI label functions as an attributional cue that shapes how viewers interpret an observed violation. For image-violating prerolls, disclosure may ease interpretation of unusual visual configurations, restoring fluency and supporting engagement; for behavior-violating prerolls, it may remove the narrative mystery that sustains curiosity. The second is product type: Hedonic products carry intrinsic affective appeal that may compete with the curiosity and elaboration associated with resolving attributional ambiguity. Utilitarian products, in contrast, may leave greater scope for these processes to contribute to user responses [32–35]. We also examine whether these moderation patterns are consistent with the attributional-ambiguity account.

This study examines how the design of AI-mediated content shapes user responses on short-video platforms, organized around the attributional-ambiguity construct. We focus on three central questions:

1. How do image-violating and behavior-violating prerolls differ in their effects on engagement and conversion, and how can attributional ambiguity explain these differences?
2. How does AI disclosure, as an attributional cue, differentially moderate user responses to image-violating versus behavior-violating prerolls?
3. How does product type (hedonic vs. utilitarian) interact with the attributional-ambiguity mechanism to shape user responses to AI-generated violations?

Using a field experiment on Chinese short-video platforms, we document four patterns interpreted through the attributional-ambiguity framework. First, conditional on exposure, behavior-violating prerolls increase sustained engagement, measured primarily by the 99% view

rate, whereas the corresponding effect of image-violating prerolls is not statistically significant. Supplementary analyses show that image-violating prerolls reduce click rate and average view duration. The sustained engagement benefit of behavior violations is consistent with our attributional-ambiguity account: Interpreting an unexpected actor–action relationship may favor an attribution to intentional AI creativity, sustaining curiosity, and attention. Second, both image-violating and behavior-violating prerolls raise conversion. The positive conversion effects are consistent with our prediction that the consequences of attributional processing can extend beyond engagement and contribute to downstream purchase behavior. Third, relative to AI conventional prerolls, AI disclosure has a positive interaction with image violations and a negative interaction with behavior violations for sustained engagement. These patterns are consistent with disclosure easing interpretation of unusual imagery while dampening the causal curiosity associated with unexpected actor–action relationships. Fourth, hedonic products weaken the effects of both violation types, consistent with hedonic crowding-out. An online survey (N = 151) provides complementary evidence consistent with the proposed attributional-ambiguity mechanism.

This study advances the IS literature on AI-mediated content production and user behavior on digital platforms in four ways. First, we introduce attributional ambiguity as the AI-specific construct that distinguishes AI-generated violations from traditional human-crafted incongruity. While expectation-violation theories explain why incongruity captures attention [18–20], they do not directly address whether an unusual AI-generated output reflects intentional creativity or generative-model error. We argue that uncertainty between these interpretations is central to understanding user responses to AI-generated content and anchor this mechanism in the algorithm-aversion and algorithm-appreciation literature [23–27]. Second, we draw on dual-stage processing to provide a cognitive foundation for our attributional-ambiguity account [28–31,36]. We explain how the structural distinction between an object's visual attributes and an actor–action relationship foregrounds perceptual coherence versus narrative interpretation, informing theoretical predictions about whether attributional ambiguity is resolved toward generative-model error or intentional AI creativity. Third, we extend the AI-disclosure literature [27,37] by showing that disclosure differentially moderates sustained engagement across violation types, consistent with its role as a cue that shapes the resolution of attributional ambiguity. Fourth, we document a hedonic crowding-out pattern in the AI-content setting, consistent with hedonic appeal competing with the curiosity and elaboration generated by attributional ambiguity, thereby extending prior work on hedonic and utilitarian consumption [32–35,38]. Collectively, our framework provides a unified mechanism linking AI-mediated content design to engagement, conversion, and disclosure effects on digital platforms.

Our findings offer two main practical implications. For advertisers, AI prerolls should be designed with the cognitive processing that the violation foregrounds in view. Behavior-violating prerolls offer a promising approach to sustaining engagement. The disclosure interactions suggest that AI labeling may ease interpretation of image violations but diminish the curiosity associated with behavior violations. Advertisers should therefore consider disclosure and creative design jointly. In addition, hedonic appeal may compete with AI-induced curiosity and elaboration, suggesting greater scope for AI-generated violations to enhance sustained engagement in utilitarian product categories. For policy makers, the heterogeneous effects of disclosure highlight the importance of considering how labeling interacts with content design. Disclosure may ease interpretation when violations foreground perceptual coherence but dampen curiosity when they invite narrative interpretation. These findings can inform the evaluation of AI-disclosure practices across creative formats.

2. Literature review

We review research on AI in advertising, algorithm aversion and attribution, and the cognitive processing of incongruity.

2.1. AI in advertising

Generative AI has rapidly become a productive input to digital advertising. Prior research highlights three themes. First, AI-generated content (AIGC) enables rapid creative iteration and large-scale dissemination; human-in-the-loop large-language-model pipelines significantly improve ad-copy generation across budget levels [39,40], and AI efficacy varies by sector [41]. Second, consumer responses to AI-generated content are task-dependent: Openness toward AI coexists with concerns about credibility and authenticity [42–45]. Third, AI disclosure has nuanced effects: It can increase transparency but reduce perceived surprise or authenticity, affecting downstream engagement [27,37,46,47]. Closely related is research on attention-capturing advertising content. Highly salient or incongruent content can break through information overload, transferring attention to the focal brand or product [21,22,48–53]. Attention-capturing prerolls operate as brief gateways that interrupt automated browsing [54–58], though misaligned or overly entertaining prerolls can hinder attention transfer [12,59,60]. Recent work suggests that AI-generated hyperpersonalized content can increase attention and curiosity [61–63].

The effectiveness of attention-capturing content also depends on the product context. Research distinguishes hedonic and utilitarian dimensions of consumption, emphasizing experiential enjoyment and functional value, respectively [32,33]. Related work examines how consumers justify hedonic versus utilitarian choices [34]. These perspectives raise a question for AI-generated advertising: whether the appeal of the advertised product complements or competes with the curiosity and elaboration elicited by unusual content.

Despite these advances, prior research has paid less attention to how different forms of AI-generated prerolls shape user responses and the cognitive mechanisms underlying these effects. Our study advances this literature by introducing attributional ambiguity—viewers' uncertainty about whether AI-generated expectation violations reflect intentional creativity or generative-model error. By distinguishing image-violating from behavior-violating prerolls, we explain how the structure of AI-generated content shapes the resolution of this ambiguity and develop theoretical predictions about engagement and conversion. We further examine how product hedonicity shapes these responses, developing a hedonic crowding-out account in which product-driven affective stimulation competes with the curiosity and elaboration associated with attributional ambiguity.

2.2. Algorithm aversion and attribution

A second stream concerns how viewers form judgments about algorithmically produced outputs. Broader research examines AI's implications for human responses and organizational capabilities [64,65]. More directly relevant to viewers' judgments, research on algorithm aversion and appreciation shows that responses depend on how people interpret algorithmic performance and capabilities. After observing the same mistake, people lose confidence more quickly in algorithmic than human forecasters [23]. Reliance on algorithms also depends on whether a task is perceived as subjective or objective and whether algorithms are believed to possess the abilities required for that task [24]. Related studies examine resistance to AI in medical decisions and responses to AI disclosure in customer interactions [26,27]. Algorithm appreciation research further shows that identifying the source as an algorithm can increase reliance on its judgments in estimation tasks [25], suggesting that algorithmic origin can also serve as a positive cue in evaluating an output. Together, these findings highlight the importance of beliefs about algorithmic performance and capabilities in shaping user responses.

Attribution research shows that people explain behavior differently depending on whether they perceive it as intentional or unintentional: Intentional behavior is explained through the actor's reasons, whereas unintentional behavior is explained through its causes [66]. We bring

this distinction to AI-generated creative content, where an unusual output can be interpreted as either a deliberate creative choice or an unintended generative-model error. Extending research on algorithm aversion following observed errors [23], our account addresses how viewers interpret an unusual output as an error or a creative choice in the first place. Our contribution is to conceptualize the uncertainty between these competing interpretations as attributional ambiguity, connecting research on responses to algorithmic errors with the attributional process through which an unusual output is interpreted.

This perspective also informs our treatment of AI disclosure. Prior research examines how identifying an output as AI-generated affects consumer responses [27,37]. We extend this question by considering how disclosure shapes the interpretation of an unusual output. Specifically, we conceptualize disclosure as an attributional cue whose consequences depend on the structure of the violation and the explanation it makes more plausible. This connects heterogeneous disclosure effects to the resolution of attributional ambiguity.

2.3. Cognitive processing of incongruity

Cognitive conflict and expectation-violation theories establish that incongruity can capture attention and reorganize downstream processing [17–22,67]. Understanding responses to different forms of incongruity also requires considering the information viewers process to make sense of the violation. Research on visual attention and narrative comprehension provides complementary foundations for our dual-stage processing perspective, which distinguishes the early registration of perceptual information from the more elaborative interpretation of depicted events [29–31,36].

For visual incongruity, research on visual attention explains how basic visual features are registered preattentively and subsequently selected and integrated into object representations [30,36]. Research on processing fluency further links the ease of processing to evaluative responses [28,68,69]. Together, these perspectives provide a basis for understanding image violations: An unusual appearance or material composition foregrounds the task of reconciling an object's visual attributes with its familiar representation, making perceptual coherence central to this interpretive task.

For behavioral incongruity, understanding the violation requires interpreting an actor's action in relation to its expected capabilities and the depicted situation. Research on mediated message processing explains how viewers allocate limited cognitive resources to encoding, storage, and retrieval [29], while research on narrative transportation highlights cognitive and affective involvement in depicted events [31]. These perspectives provide a basis for understanding why behavior violations invite narrative interpretation of an unexpected actor–action relationship, including causal questions about why the action occurs.

Building on these foundations, we use dual-stage processing to explain how the structure of a violation shapes the resolution of attributional ambiguity. Our contribution is to connect perceptual coherence and narrative interpretation to viewers' competing explanations of AI-generated incongruity as generative-model error or intentional AI creativity. These different bases for interpretation inform our theoretical predictions about how viewers resolve attributional ambiguity and how their interpretations shape engagement and conversion.

3. Theoretical framework and hypotheses

We conceptualize the effects of AI-generated prerolls on two key advertiser outcomes along the consumer response process: *engagement* and *conversion*. We distinguish two forms of expectation violation — *AI image-violating* and *AI behavior-violating* prerolls — according to the structure of the incongruity. We develop *attributional ambiguity* as the central theoretical mechanism and draw on dual-stage processing to explain how this ambiguity is resolved for each form. We further examine how their effects are moderated by *AI disclosure* and *product type*.

3.1. Attributional ambiguity as an AI-specific mechanism

Cognitive conflict and expectation-violation theories establish that moderate incongruity generates attention, curiosity, and elaboration [18–20,48,50]. These accounts, however, presume a stable attributional environment in which incongruity is read as the outcome of intentional human design. AI-generated content disrupts this assumption: When viewers encounter incongruity in an AI-produced preroll, they must evaluate not only what the violation means but also why it is there—whether it reflects intentional creative design or unintended algorithmic failure. We refer to uncertainty between these competing interpretations as attributional ambiguity.

Attributional ambiguity is distinct from — and additive to — the cognitive conflict that any incongruity produces. Whereas cognitive conflict captures the gap between expectation and observation, attributional ambiguity captures the gap between competing explanations for that observation: Intentional AI creativity versus generative-model error. The possibility of generative-model error introduces an AI-specific explanation for incongruity. This interpretive uncertainty connects our account to research on algorithm aversion and appreciation, which examines how beliefs about algorithmic performance and capabilities shape user responses [23–27].

How attributional ambiguity is resolved depends on the cognitive processing that the violation foregrounds. We conceptualize image and behavior violations according to the locus of the incongruity. Image violations concern the attributes constituting an object's visual representation: The object's appearance or material composition conflicts with expectations about what it normally looks like. Behavior violations concern the relationship between an actor and an action within an event: The action conflicts with expectations about the actor's behavior or capabilities. These properties imply different interpretive tasks. Interpreting an image violation centers on reconciling the object's appearance with its familiar representation; interpreting a behavior violation requires relating what the actor does to its capabilities and the depicted situation.

The dual-stage processing perspective provides a cognitive foundation for explaining how these different interpretive tasks shape the resolution of attributional ambiguity. Drawing on research on visual attention and narrative comprehension, we distinguish a *preattentive stage*, at which visual features are initially registered, from a *narrative stage*, at which viewers integrate information about actors and actions into a coherent account of the depicted event [28–31,36]. Image-violating prerolls foreground an object's visual configuration, with unusual visual features providing early cues to incongruity during preattentive processing. This emphasis on perceptual coherence makes the quality of the generated visual representation salient, which may favor an attribution to generative-model error, potentially eroding processing fluency and credibility [68,69]. Behavior-violating prerolls, in contrast, foreground an actor–action relationship. Understanding this relationship engages narrative processing and provides a basis for causal questions (why is the dog riding a motorcycle?), potentially favoring an attribution to intentional AI creativity that sustains curiosity and continued viewing.

In our framework, attributional ambiguity is the central theoretical mechanism. It specifies the interpretive problem — whether the oddity reflects intentional creativity or generative-model error — while dual-stage processing explains how the structure of the violation foregrounds different information for resolving that problem. Thus the dual-stage processing perspective grounds our attributional ambiguity account in the cognitive properties of the two forms of violation, informing our theoretical predictions about engagement and conversion.

3.2. Engagement and conversion

Building on the above attributional ambiguity mechanism, we expect image and behavior violations to have opposite effects on engagement (e.g., sustained viewing). For AI image-violating prerolls, an emphasis on perceptual coherence may favor an algorithmic-error attribution that reduces perceived credibility and dampens viewers' willingness to keep watching [28,68]. For AI behavior-violating prerolls, narrative-stage causal reasoning may favor an attribution to intentional AI creativity, which sustains curiosity and elaboration and encourages continued viewing [31,70].

Hypothesis 1 (*H1: Engagement*)**.** Conditional on exposure, AI behavior-violating prerolls enhance engagement, whereas AI image-violating prerolls reduce engagement.

We propose that the consequences of attributional processing can persist after viewers form an explanation for the violation. Prior research links advertising exposure and processing to subsequent responses [50,51]. We refer to the continued cognitive and motivational activation following interpretation as a *cognitive-residue* effect. Although image and behavior violations may favor different attributions, the effort involved in interpreting either form can leave lingering novelty, curiosity, or reflection that encourages subsequent product consideration. Both forms can therefore increase conversion despite their predicted differences in engagement.

Hypothesis 2 (*H2: Conversion*)**.** At the conversion stage, AI image-violating and AI behavior-violating prerolls both raise conversion by leaving residual cognitive activation that promotes goal-directed action.

3.3. Moderation by AI disclosure

AI disclosure operates as an explicit attributional cue that shapes how viewers interpret an observed violation [27,37]. Its consequences depend on the processing demands that the violation foregrounds. For image violations, identifying the AI source may make unusual visual configurations easier to interpret, improving processing fluency and supporting engagement [24,69]. For behavior violations, the same cue may reduce the narrative mystery that sustains curiosity, weakening engagement. Our cognitive-residue account further predicts that the continued activation generated during interpretation can contribute to purchase behavior, with more limited additional moderation by disclosure at the conversion stage.

Hypothesis 3 (*H3: Disclosure Effects*)**.**
AI disclosure increases engagement for image-violating prerolls but decreases engagement for behavior-violating prerolls; its effect on conversion is limited.

3.4. Product heterogeneity: Hedonic crowding-out

Product type shapes the processing goals viewers bring to an ad: Hedonic products activate affective, experiential processing, whereas utilitarian products activate analytical, goal-directed evaluation [32,33,35,38]. We argue that this processing-goal difference moderates the attributional-ambiguity mechanism in a specific way. Hedonic products already deliver intrinsic stimulation and enjoyment, so the curiosity and elaboration produced by AI-induced incongruity become partially redundant. This produces a *hedonic crowding-out effect*: Product-driven affect competes with AI-induced novelty for limited attentional resources, weakening sustained engagement and conversion [33,34]. In utilitarian contexts, by contrast, intrinsic product stimulation is weaker, so attributional processing plays a more central role in sustaining attention and transferring it toward product-relevant information.

Table 1
Summary of theoretical foundations.

Outcome	AI Image-Violating	AI Behavior-Violating	Moderation by AI Disclosure	Moderation by Hedonicity
Engagement	↓ Perceptual incoherence may favor generative-model-error attribution, reducing fluency and credibility. *[Attributional Ambiguity: Perceptual Coherence]*	↑ Narrative interpretation may favor intentional AI creativity, sustaining curiosity and attention. *[Attributional Ambiguity: Narrative Interpretation]*	↑ (Image) Eases interpretation and supports processing fluency. ↓ (Behavior) Reduces narrative mystery and curiosity. *[Attributional Cue]*	↓ Hedonic appeal competes with AI-induced curiosity and elaboration. *[Hedonic Crowding-Out]*
Conversion (Purchase)	↑ Residual cognitive activation following attributional interpretation supports subsequent product consideration and purchase. *[Cognitive Residue]*	↑ Residual cognitive activation following attributional interpretation supports subsequent product consideration and purchase. *[Cognitive Residue]*	Limited moderation Residual cognitive activation can continue to support purchase after attributional interpretation. *[Cognitive Residue]*	↓ Hedonic appeal competes with AI-induced curiosity and elaboration. *[Hedonic Crowding-Out]*

Hypothesis 4 (*H4: Product Hedonicity*)**.** Hedonic product contexts shift the effects of AI image-violating and behavior-violating prerolls on engagement and conversion in a negative direction relative to utilitarian contexts, because product-driven affective stimulation competes with the curiosity and elaboration associated with attributional ambiguity.

Table 1 summarizes the proposed relationships. Attributional ambiguity is the central mechanism, while dual-stage processing provides its cognitive foundation. The structure of the violation foregrounds different information for attributional interpretation, informing predictions about engagement. Cognitive residue connects this interpretation to subsequent product consideration and conversion. AI disclosure shapes attributional interpretation, while product hedonicity may compete with the curiosity and elaboration associated with this process.

4. Research design

We present the experimental design and the outcome variables in Sections 4.1 and 4.2.

4.1. Experimental design

We test the impact of different preroll types on consumer engagement and conversion behavior, with a focus on the moderating effect of AI disclosure and product type.

Preroll types. We study five preroll types, which are categorized as follows: (i) None: No preroll advertisement before the main video content. (ii) Conventional: A standard, non-AI-generated preroll. (iii) AI Conventional: A preroll created by AI but with content that aligns with conventional expectations (e.g., nonsurreal). (iv) AI Image-Violating: A preroll generated by AI with visually surreal or unexpected imagery (e.g., the vegetable dog). (v) AI Behavior-Violating: A preroll that includes AI-generated content where the actions or behaviors of the objects or characters defy logic (e.g., the dog riding a motorcycle).

Table 2
Experimental design.

Preroll Type	Video count	Preroll variants	Disclosure	Platforms	Days	Observations
No Preroll	12 Videos (6 products, 2 videos each)	–	–	2 Platforms	15	360 (12*2*15)
Conventional Preroll	12 Videos	4 Prerolls (2 landscape, 2 animals)	–	2 Platforms	15	1440 (12*4*2*15)
AI Conventional Preroll	12 Videos	4 Prerolls (2 landscape, 2 animals)	2 Types (disclosure or not)	2 Platforms	15	2880 (12*4*2*2*15)
AI Image-Violating Preroll	12 Videos	4 Prerolls (2 landscape, 2 animals)	2 Types (disclosure or not)	2 Platforms	15	2880 (12*4*2*2*15)
AI Behavior-Violating Preroll	12 Videos	4 Prerolls (2 landscape, 2 animals)	2 Types (disclosure or not)	2 Platforms	15	2880 (12*4*2*2*15)

AI disclosure. When AI-generated prerolls are used — including AI conventional, image-violating, and behavior-violating types — AI disclosure can be classified into two levels. (i) Disclosed: The preroll is explicitly labeled as AI-generated. (ii) Undisclosed: The preroll is not explicitly labeled as AI-generated.

Product contexts. Two product categories are tested in this experiment. (i) Utilitarian Products: These are products that fulfill basic needs, such as daily-use items, condiments, and cleaning paper products. (ii) Hedonic Products: These are products that provide sensory pleasure or novelty, such as clothing, bags, and fresh foods.

Sample and data collection. Table 2 summarizes the experimental setup. The main-ad videos are standardized short product walk-through advertisements covering six product categories. These main videos are identical across all experimental conditions, ensuring that observed differences in user responses are attributable to the preroll manipulation rather than variation in product content. In addition, all prerolls are designed under a strict stimulus equivalence framework. Specifically, all preroll conditions share the same content domains (landscape and animal scenes), identical duration, and consistent placement immediately before the main advertisement. For example, the same animal image is used as the base stimulus: It is shown as-is in the conventional condition, while in the AI conditions it is used as the input for generating (i) image-violating outputs, (ii) behavior-violating outputs, and (iii) AI-generated nonviolating variants. This ensures that differences across conditions arise solely in preroll type (conventional vs. AI-generated) and violation structure (image-violating vs. behavior-violating).

The videos were launched on major short-video platforms in China, including Douyin (the Chinese version of TikTok) and Kuaishou, which provide access to a large and active user base. All videos were released at 12:00 PM on August 15, 2025 to control for potential time-of-day effects. Data were collected from August 16, 2025 to August 30, 2025, after which traffic declined and was insufficient for further analysis. As shown in Table 2, the experiment included five preroll conditions: no preroll, conventional preroll, AI conventional preroll, AI image-violating preroll, and AI behavior-violating preroll. Each condition consisted of 12 videos covering 6 products, with identical product content. For prerolls, there were 2 landscape-related and 2 animal-related variations. For AI-generated prerolls, both disclosed and undisclosed versions were included. Across the two platforms and the 15-day data collection period, this design yielded a total of 10,440 observations, providing a rich dataset for analyzing the effects of preroll type, AI disclosure, and product category on user behavior.

4.2. Variables

We collect the following outcome variables capturing consumer behavior in terms of *engagement* and *conversion*, as shown in Table 3.

Engagement. We use the 99% view rate as the primary measure of engagement. It is defined as the proportion of views that reach 99% of the video's length. This measure captures sustained viewing through near completion and aligns with our theoretical focus on whether viewers continue watching after encountering an expectation violation. We also examine two complementary behavioral measures in supplementary analyses: click rate, defined as the ratio of product clicks to total views, and average view duration, defined as the average time spent watching the video.

Conversion. Conversion Rate: The ratio of purchases to total views, capturing the video's effectiveness in driving actual purchase behavior.

5. Experimental results

We present the summary statistics and the regression results in Sections 5.1 and 5.2.

5.1. Summary statistics

Appendix Table EC.1 reports the mean, standard deviation (SD), and number of observations (Observations) for each variable across preroll types. Fig. 2 presents the mean values of the primary outcomes, while Appendix Figure EC.1 presents the supplementary outcomes.

Engagement. AI behavior-violating prerolls have the highest 99% view rate (2.25%), followed by AI conventional (1.52%), conventional (1.44%), AI image-violating (1.30%), and none (1.10%). Pairwise comparisons indicate that AI conventional and AI behavior-violating prerolls have higher completion rates than None ($p = 0.028$ and $p = 0.015$, respectively). The difference between AI image-violating and none is not statistically significant ($p = 0.298$).

Conversion. Mean conversion rates are low across all groups. AI behavior-violating has the highest mean (0.072%), followed by AI image-violating (0.056%) and AI conventional (0.055%). The corresponding means for none and conventional are 0.002% and 0.00077%, respectively. These results suggest that AI-generated prerolls yield modestly higher conversion rates than non-AI prerolls.

We also examine click rate and average view duration as supplementary outcomes. AI image-violating prerolls have the lowest click rate (1.33%) and average view duration (1.08 s), compared with 1.81% and 2.38 s for None. The corresponding differences are significant at the 10% level for click rate ($p = 0.068$) and the 1% level for average view duration ($p < 0.001$). For AI behavior-violating prerolls, click rate (1.85%) and average view duration (2.14 s) do not differ significantly from none ($p = 0.897$ and $p = 0.549$, respectively).

Overall, the higher completion rate for behavior-violating prerolls is consistent with our theoretical prediction that interpreting an unexpected actor–action relationship may favor an attribution to intentional

Table 3
Variable descriptions.

Variable	Description
Engagement	
99% View Rate	Primary measure: The proportion of views that reach 99% of the video's length.
Click Rate	Supplementary measure: The ratio of product clicks to total views.
Average View Duration	Supplementary measure: The average time spent watching the video.
Conversion	
Conversion Rate	The ratio of purchases to total views.

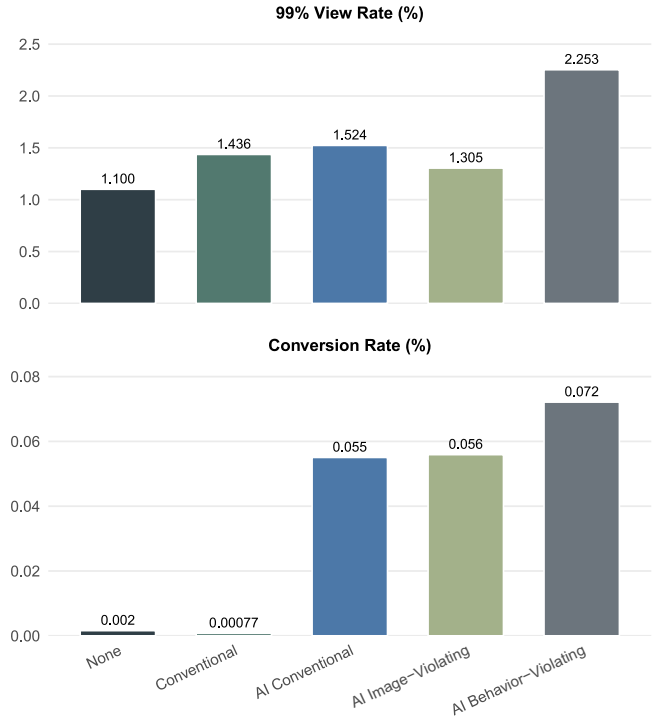


Fig. 2. Primary outcomes across preroll types.

AI creativity, sustaining curiosity, and continued viewing. For image violations, lower clicking and shorter viewing provide supplementary evidence of adverse responses, but the primary engagement measure does not support the predicted reduction in completion. The higher mean conversion rates for both violation types are directionally consistent with the conversion prediction. The following regression analyses formally examine these relationships.

5.2. Regression results

We test the effect of preroll type, the interaction effect of disclosure and the preroll type, and the product heterogeneity.

5.2.1. Effect of preroll type

We estimate the following baseline specification:

$$Y_{it} = \alpha + \beta_1 \,\text{Preroll}_i + \mathbf{X_i} + T_t + \varepsilon_{it}. \quad (1)$$

The outcome variable Y_{it} is the 99% view rate or conversion rate for video i on date t. Supplementary analyses use the same specifications for click rate and average view duration. The key explanatory variable is the preroll type, which takes one of five categories: None, Conventional, AI conventional, AI image-violating, and AI behavior-violating. $\mathbf{X_i}$ denotes the control variables, including product type (six categories) and platform type (Douyin or Kuaishou). T_t captures the time fixed effects. Finally, ε_{it} is the error term. Table 4 reports the estimation results about the effects of different preroll types on performance outcomes, with the "None preroll" as the omitted group.

Table 4
Effect of Preroll type.

Dependent Variable	99% View Rate (I)	Conversion Rate (II)
AI Conventional Preroll	0.004***	0.001*
	(0.001)	(0.000)
AI Image-Violating Preroll	0.002	0.001*
	(0.001)	(0.000)
AI Behavior-Violating Preroll	0.012***	0.001**
	(0.002)	(0.000)
Conventional Preroll	0.003**	−0.000
	(0.001)	(0.000)
Controls	Yes	Yes
Time Fixed Effect	Yes	Yes
Observations	10,440	10,440
R^2	0.112	0.011

Notes. This table tests the effect of preroll type on the outcome variables. Results are based on the full sample. The omitted reference group is *None* preroll. Robust standard errors are reported in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Engagement. Column I of Table 4 reports the 99% view rate. Relative to no preroll, AI behavior-violating increases this measure by approximately 1.2 percentage points (0.012, $p < 0.01$), indicating a higher proportion of views reaching near completion. AI conventional and conventional also have positive coefficients (0.004, $p < 0.01$; 0.003, $p < 0.05$). The coefficient for AI image-violating (0.002) is not statistically significant. Thus the results partially support H1: The predicted engagement benefit of behavior violations is supported, whereas the predicted reduction for image violations is not.

Conversion. Column II reports positive coefficients for AI conventional (0.001, $p < 0.10$), AI image-violating (0.001, $p < 0.10$), and AI behavior-violating (0.001, $p < 0.05$). These estimates are consistent with H2's directional prediction, with weaker statistical evidence for image violations. The positive image-violation coefficient, alongside its nonsignificant completion effect, indicates that the evidence for conversion and sustained engagement does not follow an identical pattern.

Together, the primary results are consistent with the prediction that interpreting an unexpected actor–action relationship can sustain curiosity and continued viewing. An unusual visual configuration, however, does not produce the predicted reduction in near-complete viewing. This qualifies the engagement prediction of our attributional ambiguity account: The evidence supports the proposed benefit of behavior violations more clearly than the proposed disadvantage of image violations. The positive conversion coefficients are consistent with our prediction that the consequences of attributional processing can extend beyond engagement and contribute to downstream purchase behavior.

We also examine click rate and average view duration as supplementary outcomes (Appendix Table EC.2, Columns I–II). AI image-violating reduces click rate (−0.005, $p < 0.05$) and average view duration (−1.293, $p < 0.01$). For AI behavior-violating, neither coefficient is statistically significant. These additional results indicate that the completion-rate benefit of behavior violations does not extend uniformly to clicking or average viewing time.

Table 5
Interaction effect of disclosure and Preroll type.

Dependent variable	99% View Rate (I)	Conversion Rate (II)
AI Image-Violating	-0.005***	0.000
Preroll	(0.001)	(0.000)
AI	0.012***	0.001
Behavior-Violating	(0.003)	(0.000)
Preroll		
Disclosure	-0.002**	0.001*
	(0.001)	(0.001)
AI Image-Violating	0.006***	−0.000
Preroll×Disclosure	(0.002)	(0.001)
AI	-0.010***	−0.001
Behavior-Violating	(0.003)	(0.001)
Preroll×Disclosure		
Controls	Yes	Yes
Time Fixed Effect	Yes	Yes
Observations	8,640	8,640
R^2	0.069	0.059

Notes. This table tests the interaction effect between preroll type and disclosure. Results are based on the sample with AI prerolls. The omitted reference group is *AI Conventional* preroll. Robust standard errors are reported in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table 6
Moderating effect of hedonic product.

Dependent Variable	99% View Rate (I)	Conversion Rate (II)
AI Conventional Preroll	0.006***	0.001*
	(0.002)	(0.001)
AI Image-Violating Preroll	0.007***	0.001*
	(0.002)	(0.001)
AI Behavior-Violating Preroll	0.016***	0.001
	(0.002)	(0.001)
Conventional Preroll	0.012***	−0.000
	(0.002)	(0.000)
Hedonic Product	0.010***	−0.000
	(0.002)	(0.000)
AI Conventional ×Hedonic	−0.003	−0.001*
	(0.003)	(0.001)
AI Image-Violating ×Hedonic	-0.010***	−0.001*
	(0.003)	(0.001)
AI Behavior-Violating ×Hedonic	-0.009**	−0.000
	(0.004)	(0.001)
Conventional ×Hedonic	-0.018***	0.000
	(0.003)	(0.000)
Controls	Yes	Yes
Time Fixed Effect	Yes	Yes
Observations	10,440	10,440
R^2	0.094	0.008

Notes. This table tests the interaction effect between preroll type and hedonic product category. Results are based on the full sample. Robust standard errors are reported in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

5.2.2. Interaction effect of disclosure and preroll type

We formally test the interaction effect between preroll type and AI disclosure

$$Y_{it} = \alpha + \beta_1 \text{Preroll}_i + \beta_2 \text{Disclosure}_i + \beta_3 \text{Preroll}_i \times \text{Disclosure}_i + \mathbf{X_i} + T_t + \varepsilon_{it}, \quad (2)$$

where Disclosure_i is a dummy variable indicating whether the AI preroll identity is disclosed, and β_3 captures the interaction effect between preroll type and disclosure. In this specification, the sample is restricted to AI prerolls only, and the omitted reference group is the AI Conventional type. Table 5 reports the estimation results regarding the interaction effect between preroll type and disclosure.

Engagement. Column I of Table 5 shows that, without disclosure, AI Image-Violating has a lower 99% view rate than AI Conventional (−0.005, $p < 0.01$), whereas AI Behavior-Violating has a higher rate (0.012, $p < 0.01$). The disclosure interaction is positive for image violations (0.006, $p < 0.01$) and negative for behavior violations (−0.010, $p < 0.01$). Disclosure therefore shifts the relative completion effect of image violations upward and that of behavior violations downward.

The point estimates illustrate the magnitude of these shifts. With disclosure, the estimated image-violation difference relative to AI Conventional changes from −0.005 to approximately 0.001, while the corresponding behavior-violation difference decreases from 0.012 to approximately 0.002. These shifts indicate that disclosure narrows the relative completion disadvantage of image violations and attenuates the relative completion advantage of behavior violations. The significant interaction terms are consistent with the directional moderation predictions in H3.

Conversion. Column II reports a positive disclosure coefficient for the AI Conventional reference category (0.001, $p < 0.10$). Neither violation-by-disclosure interaction is statistically significant. The results therefore provide no evidence that disclosure moderates conversion differently for image or behavior violations relative to AI Conventional.

The completion-rate interactions are consistent with our attributional ambiguity account. Disclosure provides an interpretive cue that may ease the evaluation of unusual visual configurations, reducing the relative completion disadvantage associated with image violations. For behavior violations, the same cue may diminish the curiosity generated by interpreting an unexpected actor–action relationship, attenuating their relative completion advantage. The evidence for these differential disclosure effects is concentrated in sustained engagement; the conversion results do not establish corresponding differences in downstream purchase responses.

We also examine click rate and average view duration as supplementary outcomes (Appendix Table EC.2, Columns III–IV). Disclosure interactions on click rate are negative for both image violations (−0.007, $p < 0.01$) and behavior violations (−0.013, $p < 0.01$). For average view duration, the interaction is positive for image violations (1.024, $p < 0.01$) and negative for behavior violations (−1.064, $p < 0.01$), matching the directions of the completion interactions.

5.2.3. Product heterogeneity

We formally test the moderating effect of the hedonic product on the preroll-type impact

$$Y_{it} = \alpha + \beta_1 \text{Preroll}_i + \beta_2 \text{Hedonic}_i + \beta_3 \text{Preroll}_i \times \text{Hedonic}_i + X_i + T_t + \varepsilon_{it}, \quad (3)$$

where Hedonic_i is a dummy variable indicating whether the product is hedonic or utilitarian, and β_3 captures the interaction effect between preroll type and hedonic product. X_i indicates the platform variable. The omitted reference group is the None type. Table 6 reports the moderating effect of hedonic product category on the relationship between preroll type and performance outcomes.

Engagement. Column I of Table 6 reports positive completion effects for AI image-violating (0.007, $p < 0.01$) and AI behavior-violating (0.016, $p < 0.01$) in the utilitarian reference context. Their interactions with hedonic product type are negative: −0.010 for image violations ($p < 0.01$) and −0.009 for behavior violations ($p < 0.05$). Thus relative to no preroll, the completion effects of both violation types are less positive in hedonic than utilitarian contexts, supporting the engagement prediction in H4. Hedonic products increase completion in the no-preroll condition (0.010, $p < 0.01$) but reduce the additional completion benefit associated with image- and behavior-violating prerolls. This pattern is consistent with the proposed hedonic crowding-out account: Product-driven affective stimulation may reduce the contribution of AI-induced curiosity to sustained viewing.

Conversion. Column II reports a negative hedonic-product interaction for AI image-violating (−0.001, $p < 0.10$), whereas the interaction for AI behavior-violating is not statistically significant. The image-violation result is consistent with our prediction that hedonic product appeal competes with the cognitive elaboration generated by attributional ambiguity, weakening its contribution to downstream purchase behavior.

The negative completion-rate interactions are consistent with the proposed hedonic crowding-out account, in which product-driven affective stimulation competes with AI-induced curiosity and elaboration. For conversion, the interaction is negative and significant for image violations but not statistically significant for behavior violations. Thus the evidence supports the engagement prediction, with more limited support for the conversion prediction.

We also examine click rate and average view duration as supplementary outcomes (Appendix Table EC.2, Columns V–VI). Hedonic-product interactions on click rate are negative for image violations (−0.017, $p < 0.01$) and behavior violations (−0.011, $p < 0.01$). In contrast, the interactions on average view duration are positive for both types (2.151 and 2.117, respectively; $p < 0.01$).

6. Robustness

We examine whether the main patterns persist across alternative specifications and sample definitions. The appendix tables report the primary outcomes—99% view rate and conversion rate—in Columns I–II and supplementary outcomes—click rate and average view duration—in Columns III–IV.

6.1. Postexposure behavior

Our estimates focus on how preroll design shapes postexposure engagement and conversion outcomes. Our primary outcomes capture sustained viewing and purchase behavior through the 99% view rate and conversion rate. Click rate and average view duration provide supplementary evidence. By construction, these measures capture user behavior after exposure rather than exposure itself. While the recommendation algorithm may affect the probability of exposure, the postexposure behavioral responses we study are driven by users' interactions with the content once exposure has occurred. All videos are further subject to the same platform recommendation algorithm regardless of preroll assignment. Therefore our identification strategy does not rely on differences in algorithmic rules across treatments; instead, it leverages variation in user responses conditional on a common exposure-generating mechanism, isolating the causal impact of preroll design on engagement and conversion outcomes.

To provide a comprehensive characterization including the exposure stage, we additionally examine effects with (i) exposure effects and robustness checks on (ii) postexposure behavioral responses.

6.1.1. Exposure (Total views)

We first examine exposure measured by total views, summarized in Appendix Table EC.3. Appendix Table EC.4 reports level, log, and Poisson specifications. Across these specifications, AI prerolls receive fewer views than the no-preroll reference condition. In the log specification, the coefficients are −0.606 for AI conventional, −0.604 for AI image-violating, and −0.574 for AI behavior-violating (all $p < 0.01$).

6.1.2. Positive views subsample

Our main analyses retain the full sample, including video-day observations with zero views, which reflect outcomes of the platform-user process rather than missing observations. For these observations, outcome rates are coded as zero. We also estimate the models on the positive-views subsample ($N = 5191$) to examine whether the main patterns persist conditional on realized viewing.

For the primary outcomes, AI behavior-violating increases the 99% view rate (0.023, $p < 0.01$), while the AI image-violating coefficient (0.004) is not statistically significant (Column I). Both image and behavior violations increase conversion rates (0.002 and 0.003, respectively; $p < 0.05$; Column II). Thus the main completion and conversion patterns persist among observations with positive views.

We also examine the supplementary outcomes (Columns III–IV). Neither violation type has a statistically significant click-rate coefficient. Image violations reduce average view duration (−1.740 s, $p < 0.01$), while behavior violations increase it (1.646 s, $p < 0.01$).

6.1.3. Exposure-controlling estimates

To further control exposure for postexposure behavior, we reestimate engagement and conversion models on the full sample while controlling for $\log(1 + \text{total views})$ (Appendix Table EC.6). This specification assesses whether the estimates are sensitive to adjustment for observed viewing volume. The main results still hold. For the primary outcomes, AI behavior-violating retains a positive completion coefficient (0.015, $p < 0.01$), while the coefficient for AI image-violating (0.002) remains nonsignificant (Column I). Both violation types have positive conversion coefficients (0.001, $p < 0.01$; Column II).

For the supplementary outcomes (Columns III–IV), image violations have a negative click-rate coefficient (−0.004, $p < 0.10$), whereas the behavior-violation coefficient is not statistically significant. Average view duration decreases for image violations (−0.469, $p < 0.05$) and increases for behavior violations (0.551, $p < 0.05$).

Across these specifications, the positive completion effect of behavior violations and the positive conversion coefficients for both types persist. These patterns are consistent with the prediction that narrative interpretation can sustain viewing and that the consequences of attributional processing can extend beyond engagement to purchase behavior.

6.2. Placebo test

Our data set is large — with 10,440 video-day observations — but the cross-section of products (6 products) and platforms (2) is small. This means that, when computing standard errors, we rely heavily on temporal variation, and persistent shocks correlated with preroll assignment could in principle bias inference. This problem is common in field experiments. For example, Ref. [71] studied judiciary task juggling but were only able to influence the behavior of six judges, and Ref. [72] tested the effects of electronic shelf labels in retail but could only study two treated stores. Ref. [73] tested the effect of ladder-design policies with only 22 origin–destination pairs.

To assess whether our results are artifacts of this small cross-section, we conduct a randomization-inference placebo test [73,74]. Specifically, we create 1000 synthetic samples by randomly reassigning preroll labels within (product clip × platform × date) cells and reestimate Eq. (1) on each synthetic sample. Fig. 3 benchmarks the realized coefficients against the distribution of placebo coefficients. If our realized estimates were artifacts of cross-correlation within the small cross-section, they would not stand out relative to the placebo distribution. Fig. 3 compares the realized coefficients with the distributions obtained from these label reassignments.

These results indicate that our findings are not artifacts of the small cross-section. For the 99% view rate, the realized behavior-violation coefficient lies far to the right of its placebo distribution, consistent with the stronger regression evidence for behavior violations. For conversion, the realized coefficients for both violation types lie toward the right tails of their respective distributions. These patterns complement the primary regression findings. The supplementary outcome distributions are retained in the lower two rows of the figure.

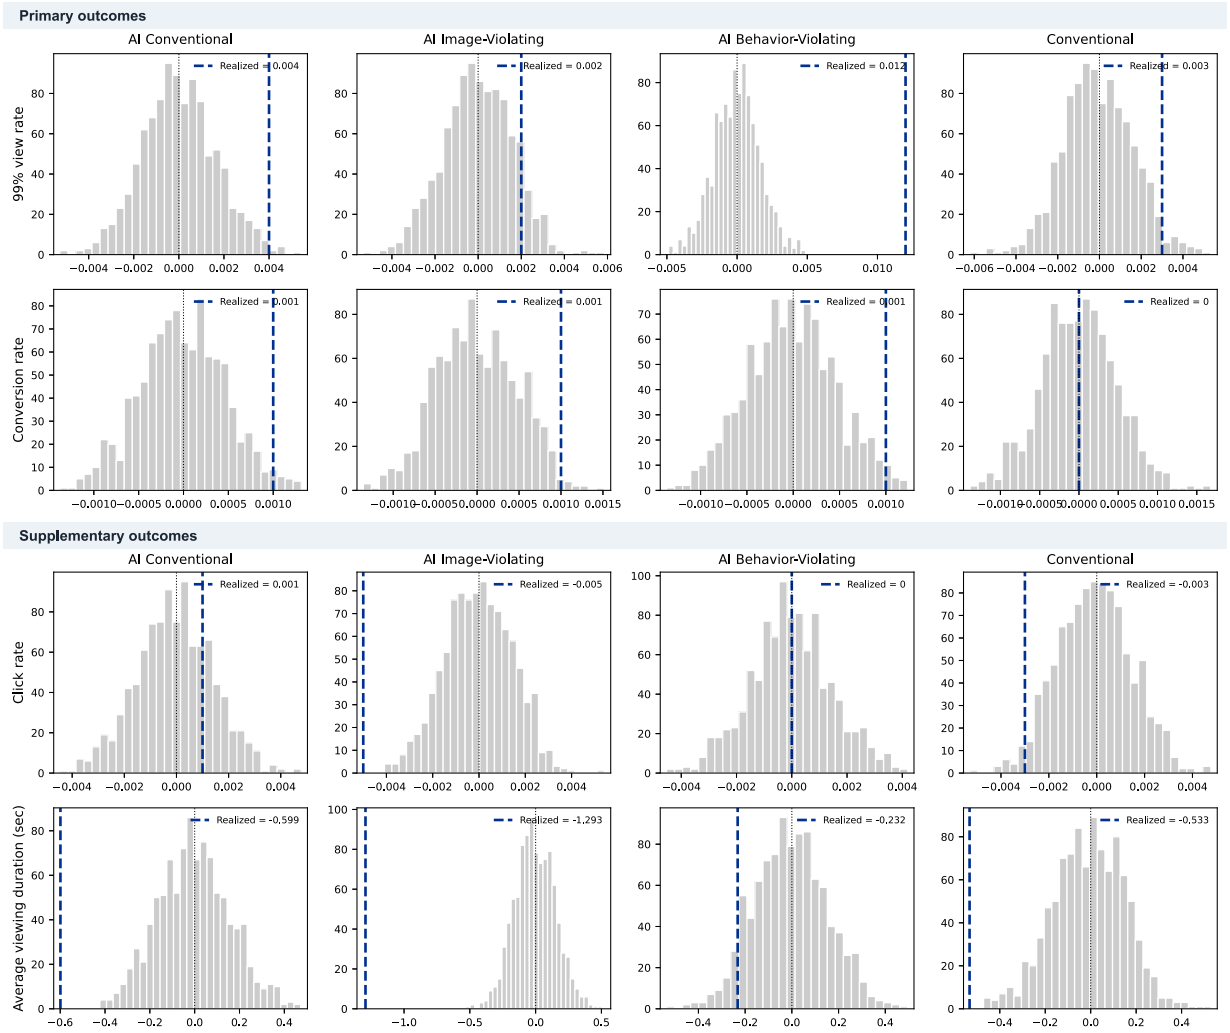


Fig. 3. Placebo coefficient distributions: Primary and supplementary outcomes.
Notes. Each panel shows 1000 placebo coefficients from preroll label reassignments within product-clip × platform × date cells. Rows 1–2 report the primary outcomes (99% view rate and conversion rate); Rows 3–4 report the supplementary outcomes (click rate and average view duration). Columns show AI conventional, AI image-violating, AI behavior-violating, and conventional prerolls. Dashed lines indicate the realized coefficient estimates.

6.3. Survey evidence

To examine whether our results reflect the cognitive mechanisms developed in Section 3, we conducted an independent online survey on a consumer panel ($N = 151$; 67 male, 84 female; mean age 34.29 years). The instrument, reproduced in Appendix EC.2, asks respondents to evaluate two short AI-generated preroll examples (Example A, an image violation; Example B, a behavior violation) on a seven-point Likert scale. The 15 items span (i) baseline interest in AI and short-video viewing, (ii) manipulation validation, (iii) attributional ambiguity, (iv) engagement intentions and purchase likelihood, (v) the moderating role of AI disclosure, and (vi) the hedonic crowding-out mechanism. Respondents report high interest in AI and frequent short-video viewing (Q1: $M = 5.78$, $SD = 0.99$, $p < 0.001$; Q2: $M = 6.01$, $SD = 1.04$, $p < 0.001$). Table EC.7 reports the full results.

Manipulation validation. The four manipulation items (Q3–Q6) support the visual–behavioral dissociation across the two stimuli. The image-violation example (Example A) scores significantly above the neutral midpoint on visual impossibility (Q3: $M = 4.47$, $p = 0.001$) but significantly below the midpoint on behavioral impossibility (Q5: $M = 3.73$, $p = 0.050$). The behavior-violation example (Example B) shows the mirror pattern, scoring significantly above the midpoint on behavioral impossibility (Q6: $M = 4.53$, $p < 0.001$) and significantly below it on visual impossibility (Q4: $M = 3.67$, $p = 0.020$). These responses support the distinction between incongruity in an object's visual representation and incongruity in an actor–action relationship.

Attributional ambiguity. Items Q7–Q8 examine attributional ambiguity and the proposed attributional tendencies. Respondents endorse the statement that the image-violation example is ambiguous between intentional design and an AI error, and that they lean toward an AI-error interpretation (Q7: $M = 4.70$, $p < 0.001$). For the behavior-violation example they likewise report ambiguity on the same dimension, but lean toward an intentional-AI-creativity interpretation (Q8: $M = 4.53$, $p < 0.001$). These responses are consistent with our prediction that the structure of a violation informs the resolution of attributional ambiguity: Image violations foreground perceptual coherence, whereas behavior violations foreground the interpretation of an actor–action relationship.

Attributional interpretation and engagement. Respondents associate the AI-error interpretation of the image-violation example with reduced willingness to engage (Q9: $M = 4.48$, $p < 0.001$) and the AI-creativity interpretation of the behavior-violation example with greater curiosity and willingness to engage (Q10: $M = 4.60$, $p < 0.001$). These responses are consistent with H1's proposed connection between attributional

interpretation and engagement: An error interpretation discourages engagement, whereas a creative interpretation sustains curiosity and attention. The behavior-violation response complements the higher 99% view rate observed in the field experiment. For image violations, the reported reduction in engagement intentions corresponds to the lower click rate and average view duration in the supplementary field analyses, although the effect on the primary engagement measure is not statistically significant. Together, these findings connect viewers' reported attributional interpretations to the engagement patterns observed across the two violation types.

Cognitive residue and conversion. Q11–Q12 examine how the cognitive activity elicited by the preroll carries forward to product consideration and purchase intentions. Respondents report that lingering novelty associated with the image-violation example increases their willingness to purchase the advertised product (Q11: $M = 4.55$, $p < 0.001$). For the behavior-violation example, continued curiosity and reflection increase this willingness (Q12: $M = 4.83$, $p < 0.001$). We also conducted follow-up interviews with two participants to explore how their interpretations of the prerolls carried forward to subsequent product consideration. Appendix Table EC.9 presents selected excerpts from these interviews.

These responses support the cognitive-residue account underlying H2: The consequences of attributional processing extend beyond interpreting the preroll and continue to motivate interest in the advertised product. Although the two violation types invite different attributional interpretations, lingering novelty for image violations and continued curiosity and reflection for behavior violations provide routes through which this processing can contribute to purchase intentions. This evidence complements the positive conversion effects of both violation types in the field experiment, connecting continued cognitive and motivational activity reported in the survey with observed purchase behavior.

AI disclosure and attributional interpretation. Respondents report that AI disclosure makes the image-violation example easier to interpret and accept (Q13: $M = 4.72$, $p < 0.001$), while reducing the mystery and curiosity associated with the behavior-violation example (Q14: $M = 4.60$, $p < 0.001$). These responses are consistent with H3's proposed role of disclosure as an attributional cue: Identifying the AI source may ease interpretation of unusual visual configurations but diminish the causal curiosity elicited by unexpected actor–action relationships. This pattern complements the positive image-violation and negative behavior-violation disclosure interactions on the 99% view rate in the field experiment. The survey thus provides an interpretive basis for the contrasting disclosure interactions observed for sustained engagement.

Product hedonicity and crowding-out. Respondents report that the novelty and curiosity triggered by AI-generated violations compete with the intrinsic appeal of hedonic products and weaken product evaluation (Q15: $M = 4.83$, $p < 0.001$). This response is consistent with H4's hedonic crowding-out account, in which product-driven affective stimulation competes with the curiosity and elaboration associated with attributional ambiguity. It complements the negative hedonic-product interactions on the 99% view rate for both violation types in the field experiment. Together, these findings suggest that the contribution of AI-generated incongruity to sustained engagement depends on the appeal of the advertised product.

Overall, the survey complements the field experiment by connecting observed behavioral patterns with viewers' reported interpretations and motivational responses. The engagement items link error and creativity interpretations to willingness to engage, while the conversion items connect lingering novelty, curiosity, and reflection to subsequent purchase intentions. The disclosure and product-hedonicity items further inform how attributional cues and product appeal shape these responses. Together, the two sources provide convergent evidence consistent with the attributional-ambiguity framework and its proposed implications for engagement and conversion.

7. Conclusion and implications

Using a field experiment on major Chinese short-video platforms, this study examines how different types of AI-generated prerolls shape user engagement and conversion behaviors. We use the 99% view rate as the primary measure of engagement and examine click rate and average view duration as supplementary outcomes. Table 7 summarizes the primary estimation results together with their theoretical interpretations.

For engagement, the effects differ across preroll types. AI behavior-violating prerolls significantly increase the 99% view rate relative to the no-preroll baseline, whereas the corresponding effect of AI image-violating prerolls is not statistically significant. Supplementary analyses show that AI image-violating prerolls significantly reduce click rate and average view duration. These patterns provide partial support for our engagement predictions and are consistent with the attributional-ambiguity framework. Drawing on dual-stage processing as its cognitive foundation, our account proposes that image violations foreground the perceptual coherence of an object's appearance, potentially favoring a generative-model-error interpretation that reduces fluency and credibility. In contrast, behavior violations foreground the interpretation of an actor–action relationship, potentially favoring intentional AI creativity and sustaining curiosity and attention. The higher 99% view rate for behavior violations is consistent with this predicted advantage in sustained engagement.

For conversion, both image-violating and behavior-violating prerolls generate positive effects on conversion rate. These findings are consistent with our prediction that the consequences of attributional processing can extend beyond engagement and contribute to downstream purchase behavior. Our account proposes that the cognitive effort involved in resolving attributional ambiguity leaves residual cognitive activation that carries forward to goal-directed behavior.

Moderating factors further shape these outcomes. Relative to AI Conventional prerolls, AI disclosure has a positive interaction with image violations and a negative interaction with behavior violations for the 99% view rate, while neither interaction is statistically significant for conversion. These results are consistent with disclosure operating as an attributional cue that changes how viewers interpret AI-generated incongruity: It may ease interpretation of unusual visual configurations while diminishing the causal curiosity associated with unexpected actor–action relationships. In addition, hedonic products weaken the completion-rate effects of both image-violating and behavior-violating prerolls. This pattern is consistent with the proposed hedonic crowding-out account, in which intrinsic product enjoyment competes with AI-induced curiosity and elaboration. For conversion, the negative hedonic-product interaction is marginally significant for image violations but not statistically significant for behavior violations, providing more limited support for the predicted moderation.

Overall, AI-generated prerolls exhibit heterogeneous effects across engagement and conversion behaviors. The findings suggest that the effectiveness of AI-mediated content depends not only on whether content violates expectations but also on how viewers cognitively interpret the source and meaning of those violations within platform-based AI advertising environments.

7.1. Theoretical implications

Our findings yield four interrelated contributions to the IS literature on AI-mediated content production and user behavior on digital platforms, organized around a single mechanism — attributional ambiguity — whose direction of resolution depends on the cognitive processing that an AI-generated violation foregrounds.

First, we introduce attributional ambiguity as the AI-specific construct that distinguishes AI-generated violations from human-crafted incongruity. Expectation-violation theories explain why incongruity captures attention [18–20], but they implicitly assume an intentional

Table 7
Summary of primary estimation results.

Outcome	AI Image-violating	AI Behavior-violating
Main Effects		
Engagement (99% View Rate)	NS Effect not statistically significant	+ Consistent with narrative interpretation sustaining curiosity and attention
Conversion Rate	+ Consistent with residual cognitive activation contributing to purchase behavior	+ Consistent with residual cognitive activation contributing to purchase behavior
Moderation: AI Disclosure		
Engagement (99% View Rate)	+ Consistent with disclosure easing interpretation and restoring fluency	– Consistent with disclosure reducing narrative mystery and curiosity
Conversion Rate	NS Interaction not statistically significant	NS Interaction not statistically significant
Moderation: Hedonic Product		
Engagement (99% View Rate)	– Consistent with hedonic appeal competing with AI-induced curiosity and elaboration	– Consistent with hedonic appeal competing with AI-induced curiosity and elaboration
Conversion Rate	– Marginal support for hedonic crowding-out in conversion	NS Interaction not statistically significant

Notes: + and – indicate positive and negative coefficients significant at the 10% level or better; NS indicates a coefficient that is not statistically significant. Main effects are relative to the no-preroll baseline. AI-disclosure moderation summarizes interaction coefficients relative to AI Conventional prerolls. Hedonic-product moderation summarizes interaction coefficients relative to the no-preroll reference. Theoretical interpretations refer to the proposed attributional-ambiguity account.

human designer and therefore do not directly address whether an unusual output reflects intentional creativity or generative-model error. By anchoring the construct in the algorithm-aversion and algorithm-appreciation literature [23–27], we identify this interpretive question as the central mechanism that differentiates AI-generated from human-crafted creative material. This perspective explains how viewers interpret an unusual output as an error or a creative choice in the first place, informing theoretical predictions about engagement and conversion.

Second, we draw on dual-stage processing to provide a cognitive foundation for explaining how the structure of a violation shapes the resolution of attributional ambiguity [28–31,36]. Image violations concern the attributes constituting an object's visual representation, whereas behavior violations concern the relationship between an actor and an action within an event. These structural properties foreground different processing demands. Unusual visual features provide early cues to image incongruity, directing interpretation toward the perceptual coherence of the object's appearance. Behavior violations invite narrative interpretation of the actor–action relationship and causal questions about the depicted event. These different bases for interpretation may favor generative-model-error and intentional-creativity attributions, respectively. Dual-stage processing thus supports our central attributional-ambiguity mechanism by explaining why the two forms of violation may lead viewers toward different interpretations.

Third, we extend the AI-disclosure literature [27,37] by deriving heterogeneous moderation predictions from the same construct: Disclosure provides an attributional cue, and its consequences depend on how viewers interpret the violation. For image violations, disclosure may restore fluency by making unusual visual configurations easier to interpret; for behavior violations, it may remove the narrative mystery that sustains curiosity. The opposing disclosure interactions for the 99% view rate are consistent with these predictions. The heterogeneous disclosure effect is therefore a direct empirical implication of the attributional-ambiguity mechanism, with support concentrated in sustained engagement.

Fourth, we document a hedonic crowding-out pattern in the AI-content setting. The negative hedonic-product interactions for the 99% view rate are consistent with hedonic appeal competing with the curiosity and elaboration generated by attributional ambiguity. This extends prior work on hedonic-versus-utilitarian processing and goal congruence [32–35,38] by explaining how product appeal may alter the contribution of AI-generated incongruity to sustained engagement. The more limited conversion evidence suggests that this crowding-out pattern is most evident in near-complete viewing.

Together, these four contributions provide a unified framework linking AI-mediated content design to user engagement, conversion behavior, and disclosure effects on digital platforms through the mechanism of attributional ambiguity.

7.2. Practical implications

The findings provide concrete guidance for advertisers and policy makers.

For advertisers operating on short-video platforms, AI prerolls should be designed with the cognitive processing that the violation foregrounds firmly in view. Behavior-violating creatives offer a promising approach when the objective is sustained viewing, as reflected in their higher 99% view rate. Image-violating creatives require attention to how viewers interpret unusual visual configurations: Their completion-rate effect is not statistically significant, while supplementary results show lower click rates and shorter average viewing duration. The disclosure interactions suggest that explicit AI labeling may ease interpretation and support sustained viewing for image violations but diminish the curiosity associated with behavior violations. Advertisers should therefore consider disclosure and creative design jointly when developing AI-generated prerolls. Hedonic affective appeal may compete with the same curiosity channel, so the incremental completion benefits of AI-generated violations are stronger in utilitarian product categories. Both violation types also show positive conversion effects, supporting the evaluation of sustained viewing and purchase outcomes as distinct advertising objectives.

For policy makers, the heterogeneous moderation of AI disclosure across violation types highlights the importance of considering how labeling interacts with content design. Disclosure may ease interpretation where a violation foregrounds perceptual coherence but dampen curiosity where it invites narrative interpretation. These findings can inform the evaluation of disclosure practices by showing how the same AI label may have different consequences for sustained viewing across creative formats.

Overall, our work provides a framework for understanding user responses to AI-generated video content in digital environments. By linking AI-generated violations, user responses, and disclosure practices through the mechanism of attributional ambiguity, we provide both

theoretical insight and actionable guidance for AI-mediated creative design and governance.

CRediT authorship contribution statement

Shichen Zhang: Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization, Writing – review & editing. **Liang Peng:** Writing – original draft, Methodology, Investigation, Formal analysis, Data curation. **Yinliang (Ricky) Tan:** Writing – review & editing, Investigation, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors thank the anonymous review team for their insightful comments and constructive suggestions, which significantly improved the manuscript. Shichen Zhang acknowledges support from the National Natural Science Foundation of China (grants 72571191 and 72201140). Liang Peng and Yinliang (Ricky) Tan acknowledge support from the Doctor of Business Administration (DBA) program at China Europe International Business School (CEIBS).

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.im.2026.104465.

References

[1] Joseph Ruff, Information overload: Causes, symptoms and solutions, Harv. Grad. Sch. Educ. (2002) 1–13.
[2] Nick Bostrom, Information hazards: A typology of potential harms from knowledge, Rev. Contemp. Philos. (10) (2011) 44–79.
[3] Thomas H. Davenport, John C. Beck, The attention economy, Ubiquity 2001 (May) (2001) 1–es.
[4] Georg Franck, The economy of attention, J. Sociol. 55 (1) (2019) 8–19.
[5] Reshaa F. Alruwaili, Scroll immersion and short-form video use: Predictors of attention, memory, and fatigue among saudi social media users, Acta Psychol. 260 (2025) 105674.
[6] Francesco Chiossi, Luke Haliburton, Changkun Ou, Andreas Martin Butz, Albrecht Schmidt, Short-form videos degrade our capacity to retain intentions: Effect of context switching on prospective memory, in: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, 2023, pp. 1–15.
[7] McKinsey & Company, The economic potential of generative AI: The next productivity frontier, 2023, https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier. (Accessed 14 May 2026).
[8] Xiahua Wei, Naveen Kumar, Han Zhang, Addressing bias in generative AI: Challenges and research opportunities in information management, Inf. Manag. 62 (2) (2025) 104103.
[9] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, Björn Ommer, High-resolution image synthesis with latent diffusion models, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 10684–10695.
[10] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, Mark Chen, Hierarchical text-conditional image generation with clip latents, 2022, arXiv preprint arXiv:2204.06125.
[11] Melissa Russell, AI will shape the future of marketing, 2025, https://professional.dce.harvard.edu/blog/ai-will-shape-the-future-of-marketing/#The-Current-State-of-AI-in-Marketing. (Accessed 14 May 2026).
[12] C. Campbell, F. Mattison Thompson, P.E. Grimm, K. Robson, Understanding why consumers don't skip pre-roll video ads, J. Advert. 46 (3) (2017) 411–423.
[13] McKinsey & Company, The state of AI in early 2024: Gen AI adoption spikes and starts to generate value, 2024, https://www.mckinsey.com/capabilities/quantumblack/%20our-insights/the-state-of-ai-2024. (Accessed 14 May 2026).
[14] TikTok, Introducing TikTok Symphony Your AI-powered content partner, 2024, TikTok Newsroom announcement, TikTok World 2024. https://ads.tiktok.com/business/creativecenter/tools/tiktok-symphony/pc/en. (Accessed 14 May 2026).
[15] Meta For Business, Generative AI features in Meta Advantage+, 2024, https://www.facebook.com/business/ (Accessed 14 May 2026).
[16] Paola Cillo, Gaia Rubera, Generative AI in innovation and marketing processes: A roadmap of research opportunities, J. Acad. Mark. Sci. 53 (3) (2025) 684–701.
[17] Eun-Ju Lee, David W. Schumann, Explaining the special case of incongruity in advertising: Combining classic theoretical approaches, Mark. Theory 4 (1–2) (2004) 59–90.
[18] Susan E. Heckler, Terry L. Childers, The role of expectancy and relevancy in memory for verbal and visual information: what is incongruency? J. Consum. Res. 18 (4) (1992) 475–492.
[19] George Mandler, The structure of value: Accounting for taste, in: Affect and Cognition, Psychology Press, 2014, pp. 3–36.
[20] Elliot Aronson, The theory of cognitive dissonance: A current perspective, in: Advances in Experimental Social Psychology, vol. 4, Elsevier, 1969, pp. 1–34.
[21] Elizabeth C. Hirschman, The effect of verbal and pictorial advertising stimuli on aesthetic, utilitarian and familiarity perceptions, J. Advert. 15 (2) (1986) 27–34.
[22] Angela Y. Lee, Aparna A. Labroo, The effect of conceptual and perceptual fluency on brand evaluation, J. Mark. Res. 41 (2) (2004) 151–165.
[23] Berkeley J. Dietvorst, Joseph P. Simmons, Cade Massey, Algorithm aversion: People erroneously avoid algorithms after seeing them err, J. Exp. Psychol. [Gen.] 144 (1) (2015) 114–126.
[24] N. Castelo, M.W. Bos, D.R. Lehmann, Task-dependent algorithm aversion, J. Mark. Res. 56 (5) (2019) 809–825.
[25] Jennifer M. Logg, Julia A. Minson, Don A. Moore, Algorithm appreciation: People prefer algorithmic to human judgment, Organ. Behav. Hum. Decis. Process. 151 (2019) 90–103.
[26] Chiara Longoni, Andrea Bonezzi, Carey K. Morewedge, Resistance to medical artificial intelligence, J. Consum. Res. 46 (4) (2019) 629–650.
[27] Xueming Luo, Siliang Tong, Zheng Fang, Zhe Qu, Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases, Mark. Sci. 38 (6) (2019) 937–947.
[28] Piotr Winkielman, Norbert Schwarz, Tetra Fazendeiro, Rolf Reber, et al., The hedonic marking of processing fluency: Implications for evaluative judgment, Psychol. Eval.: Affect. Process. Cogn. Emot. 189 (2003) 217.
[29] Annie Lang, The limited capacity model of mediated message processing, J. Commun. 50 (1) (2000) 46–70.
[30] Anne M. Treisman, Garry Gelade, A feature-integration theory of attention, Cogn. Psychol. 12 (1) (1980) 97–136.
[31] Melanie C. Green, Timothy C. Brock, The role of transportation in the persuasiveness of public narratives, J. Pers. Soc. Psychol. 79 (5) (2000) 701–721.
[32] Kevin E. Voss, Eric R. Spangenberg, Bianca Grohmann, Measuring the hedonic and utilitarian dimensions of consumer attitude, J. Mark. Res. 40 (3) (2003) 310–320.
[33] Barry J. Babin, William R. Darden, Mitch Griffin, Work and/or fun: Measuring hedonic and utilitarian shopping value, J. Consum. Res. 20 (4) (1994) 644–656.
[34] Erica M. Okada, Justification effects on consumer choice of hedonic and utilitarian goods, J. Mark. Res. 42 (1) (2005) 43–57.
[35] Joseph Cesario, Heidi Grant, E. Tory Higgins, Regulatory fit and persuasion: Transfer from "feeling right.", J. Pers. Soc. Psychol. 86 (3) (2004) 388.
[36] Jeremy M. Wolfe, Guided search 2.0: A revised model of visual search, Psychon. Bull. Rev. 1 (2) (1994) 202–238.
[37] L. Wu, N.A. Dodoo, T.J. Wen, Disclosing AI's involvement in advertising to consumers: A task-dependent perspective, J. Advert. 54 (1) (2025) 20–38.
[38] E. Tory Higgins, Making a good decision: Value from fit, Am. Psychol. 55 (11) (2000) 1217–1230.
[39] X. Yang, R. Zhan, D.F. Wong, J. Wu, L.S. Chao, Human-in-the-loop machine translation with large language model, 2023, arXiv preprint arXiv:2310.08908.
[40] M. Chui, J. Manyika, M. Miremadi, N. Henke, R. Chung, P. Nel, S. Malhotra, Notes from the AI Frontier: Insights from Hundreds of Use Cases, Technical Report, 2 (267), McKinsey Global Institute, 2018, pp. 1–31.
[41] M. Wedel, P.K. Kannan, Marketing analytics for data-rich environments, J. Mark. 80 (6) (2016) 97–121.
[42] J. Huh, M.R. Nelson, C.A. Russell, ChatGPT, AI advertising, and advertising research and education, J. Advert. 52 (4) (2023) 477–482.
[43] W.H. Kunz, J. Wirtz, Corporate digital responsibility (CDR) in the age of AI: implications for interactive marketing, J. Res. Interact. Mark. 18 (1) (2024) 31–37.
[44] J. Paul, A. Ueno, C. Dennis, ChatGPT and consumers: Benefits, pitfalls and future research agenda, Int. J. Consum. Stud. 47 (4) (2023) 1213–1225.
[45] Aida Sanatizadeh, Yingda Lu, Keran Zhao, Yuheng Hu, Engagement or entanglement? The dual impact of generative artificial intelligence in online knowledge exchange platforms, Inf. Manag. (2025) 104178.
[46] Z. Chen, J. Chan, Large language model in creative work: The role of collaboration modality and user expertise, Manag. Sci. 70 (12) (2024) 9101–9117.
[47] L. Arango, S.P. Singaraju, O. Niininen, Consumer responses to AI-generated charitable giving ads, J. Advert. 52 (4) (2023) 486–503.
[48] R. Pieters, M. Wedel, Attention capture and transfer in advertising: Brand, pictorial, and text-size effects, J. Mark. 68 (2) (2004) 36–50.

[49] Ralf Van der Lans, Rik Pieters, Michel Wedel, Research note—Competitive brand salience, Mark. Sci. 27 (5) (2008) 922–931.
[50] S. Shapiro, D.J. MacInnis, S.E. Heckler, The effects of incidental ad exposure on the formation of consideration sets, J. Consum. Res. 24 (1) (1997) 94–104.
[51] D.A. Schweidel, W.W. Moe, Listening in on social media: A joint model of sentiment and venue format choice, J. Mark. Res. 51 (4) (2014) 387–402.
[52] Jifeng Luo, Sulin Ba, Han Zhang, The effectiveness of online shopping characteristics and well-designed websites on satisfaction, Mis Q. (2012) 1131–1144.
[53] Jingchuan Pu, Young Kwark, Sang Pil Han, Qiang Ye, Bin Gu, Uncertainty reduction vs. reciprocity: Understanding the effect of a platform-initiated reviewer incentive program on regular ratings, Inf. Syst. Res. 35 (3) (2024) 1363–1381.
[54] D. Grewal, A.L. Roggeveen, Understanding retail experiences and customer journey management, J. Retail. 96 (1) (2020) 3–8.
[55] J. Alé-Chilet, S. Moshary, Beyond consumer switching: Supply responses to food packaging and advertising regulations, Mark. Sci. 41 (2) (2022) 243–270.
[56] A. Lambrecht, C. Tucker, When does retargeting work? Information specificity in online advertising, J. Mark. Res. 50 (5) (2013) 561–576.
[57] A. Kirmani, Advertising repetition as a signal of quality: If it's advertised so much, something must be wrong, J. Advert. 26 (3) (1997) 77–86.
[58] T. Teixeira, M. Wedel, R. Pieters, Moment-to-moment optimal branding in TV commercials: Preventing avoidance by pulsing, Mark. Sci. 29 (5) (2010) 783–804.
[59] J. Newton, F.J. Newton, J. Wong, Y. Tsarenko, Effect of preroll ad choice on brand attitude and attitude toward the viewing experience, Eur. J. Mark. 59 (4) (2025) 999–1024.
[60] E.A. De Leon, B.R. Ma, D.R. Virtusio, Advertising appeals influencing the brand engagement of gen Y and Z in terms of social media pre-roll advertisements, J. Bus. Manag. Stud. 4 (2) (2022) 163.
[61] V. Kumar, S. Gupta, Conceptualizing the evolution and future of advertising, J. Advert. 45 (3) (2016) 302–317.
[62] A. LeBrun, The risks of AI-generated, hyper-personalized digital advertisements, Philos. Technol. 38 (3) (2025) 1–19.
[63] A. Kapoor, M. Kumar, Frontiers: Generative AI and personalized video advertisements, Mark. Sci. (2025) Forthcoming.
[64] Xusen Cheng, Lu Gao, Xin Robert Luo, From emotion to reflection: Leveraging EmotionPrompt strategy to empower self-determination in decision-making with generative artificial intelligence, Inf. Manag. (2025) 104194.
[65] David T.W. Wong, Eric W.T. Ngai, Impact of artificial intelligence (AI) on operational performance: The role of dynamic capabilities, Inf. Manag. 62 (6) (2025) 104162.
[66] Bertram F. Malle, How people explain behavior: A new theoretical framework, Pers. Soc. Psychol. Rev. 3 (1) (1999) 23–48.
[67] T.L. Powers, E.P. Jack, The influence of cognitive dissonance on retail product returns, Psychol. Mark. 30 (8) (2013) 724–735.
[68] Rolf Reber, Norbert Schwarz, Piotr Winkielman, Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? Pers. Soc. Psychol. Rev. 8 (4) (2004) 364–382.
[69] Adam L. Alter, Daniel M. Oppenheimer, Uniting the tribes of fluency to form a metacognitive nation, Pers. Soc. Psychol. Rev. 13 (3) (2009) 219–235.
[70] Mark J. Gruber, Ben D. Gelman, Charan Ranganath, States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit, Neuron 84 (2) (2014) 486–496.
[71] Robert L Bray, Decio Coviello, Andrea Ichino, Nicola Persico, Multitasking, multiarmed bandits, and the Italian judiciary, Manuf. Serv. Oper. Manag. 18 (4) (2016) 545–558.
[72] Ioannis Stamatopoulos, Achal Bassamboo, Antonio Moreno, The effects of menu costs on supply chain efficiency: Evidence from adoption of the electronic shelf label technology, Technical Report, The University of Texas at Austin, Austin, 2017, Working paper.
[73] Maxime C Cohen, Alexandre Jacquillat, Juan Camilo Serpa, Michael Benborhoum, Managing airfares under competition: Insights from a field experiment, Manag. Sci. 69 (10) (2023) 6076–6108.
[74] Susan Athey, Guido W. Imbens, The econometrics of randomized experiments, in: Handbook of Economic Field Experiments, vol. 1, Elsevier, 2017, pp. 73–140.

Shichen Zhang is an associate professor in the College of Management and Economics at Tianjin University. Her research interests focus on online platform, AI operations, and field experiments. She has published articles in journals such as Management Science, Manufacturing & Service Operations Management, Production & Operations Management, Decision Sciences, European Journal of Operational Research, and Omega-The International Journal of Management Science.

Liang Peng was previously the CEO of Youju Interactive Beijing Technology Co. Ltd.

Yinliang (Ricky) Tan is a Full Professor of Decision Sciences and Management Information Systems at China Europe International Business School (CEIBS). His research lies at the intersection of AI economics and AI management. Professor Tan has published articles in journals such as Management Science, MIS Quarterly, Information Systems Research, Manufacturing & Service Operations Management, Production & Operations Management, and Human Relations.