Journal of
Theoretical and Applied Electronic Commerce Research

MDPI

Article

Human Spokesperson Versus Artificial Intelligence Avatar and Authorship Disclosure in Digital Advertising: Effects on Performance Metrics

Edwin Arango-Espinal [1,*], Augusto Rodríguez Orejuela [2] and Paula Andrea López-Herrera [3]

[1] Development in Specialty Coffees Group (GiiDCE), System of Regionalization, Universidad del Valle, Caicedonia 762540, Valle del Cauca, Colombia
[2] Marketing Research Group, Facultad de Ciencias de la Administración, Universidad del Valle, Cali 25360, Valle del Cauca, Colombia; augusto.rodriguez@correounivalle.edu.co
[3] SINERGIA Research Group, Institución Universitaria de Roldanillo (UNINTEP), Roldanillo 761550, Valle del Cauca, Colombia; paula.lopez.herrera@correounivalle.edu.co
* Correspondence: edwin.arango@correounivalle.edu.co

Abstract

Virtual influencers generated by artificial intelligence have emerged as a strategic tool in digital advertising, although evidence regarding their actual effectiveness in market campaigns remains limited to laboratory experiments based on self-report measures. This study examines the effect of advertising spokesperson type (human versus AI-generated avatar) and authorship disclosure (stating or not stating the origin of the advertisement) on the performance metrics of an actual campaign on the Meta platform. Through a field experiment with a 2 × 2 factorial design (spokesperson type × authorship disclosure), stratified by gender and age group, 384 advertisements were placed with a fixed budget on Facebook, and performance metrics (reach, video plays, clicks, and messages) were modeled using generalized structural equation modeling (GSEM) with a negative binomial distribution. The results reveal a trade-off: the AI avatar is associated with 25% more video plays, but with 13% less reach and 38% fewer clicks, and with no significant difference in messages. Because impressions were allocated by the platform's delivery algorithm, these estimates are read as performance differences under standard algorithmic delivery rather than as effects of full random assignment of users to conditions. Authorship disclosure did not show a statistically significant effect on any of the evaluated metrics, nor did its interaction with spokesperson type, suggesting that stating the origin of the advertisement can be implemented without a detectable cost to performance, although for messages the evidence remains inconclusive. The originality of this study lies in providing empirical evidence based on real behavioral data from an advertising campaign, overcoming the limitations of experiments relying on self-report measures that predominate in the literature.

Keywords: artificial intelligence; virtual influencers; avatars; advertising performance; authorship disclosure; digital marketing; field experiment

Check for updates

Academic Editor: Ting Chi

Received: 16 August 2026
Revised: 17 September 2026
Accepted: 18 September 2026
Published: 23 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Virtual influencers generated by artificial intelligence attract more visual attention than human spokespersons, yet this advantage does not translate into more favorable evaluations or stronger purchase intentions [1,2]. This paradox—capturing the gaze without mobilizing action—encapsulates one of the central dilemmas facing the advertising

industry in its transition toward the automation of persuasive communication [3]. Artificial intelligence (AI) has established itself as a transformative technology in the global economy, fundamentally altering business strategies, productivity, and the ways organizations communicate with their audiences [4]. In the realm of digital marketing, this transformation manifests in unprecedented levels of personalization and efficiency spanning from predictive analysis of consumer behavior to the automated generation of synthetic content [5]. One of the most visible expressions of this technological revolution is the emergence of virtual influencers (VIs), synthetic entities whose market is experiencing exponential growth driven by strategic advantages such as total control over the message, mitigation of reputational risks, and uninterrupted availability for global campaigns [6]; see Section 2.1.

However, the effectiveness of AI avatars as advertising spokespersons is far from a settled question. The literature documents a fundamental tension between the strategic advantages of these synthetic entities and the psychological barriers of consumers: human spokespersons continue to be perceived as more credible and trustworthy than their artificial counterparts [2,7,8], virtual influencers attract more initial visual attention but generate inferior attitudes and purchase intentions [1], and disclosure of the artificial nature of the spokesperson can activate persuasion resistance mechanisms that erode credibility and purchase intention [9–11]. Nevertheless, the majority of these findings stem from laboratory experiments based on self-report measures, such as purchase intention, attitude toward the advertisement, or perceived credibility, which, while valuable for understanding the underlying psychological mechanisms, do not capture the actual behavior of consumers in their natural media consumption environment.

This reliance on self-reported measures constitutes a significant limitation of the existing body of knowledge. There is a notable scarcity of research evaluating the impact of AI-generated spokespersons on the real performance metrics of advertising campaigns on digital platforms [12]. The aggregate behavioral data reported by advertising platforms—reach, plays, clicks, conversions—offer an ecologically valid measure of consumer response that is not subject to social desirability biases or the gap between stated intention and actual behavior. Similarly, the question of whether AI authorship disclosure affects campaign performance under real market conditions remains without a satisfactory empirical answer. The available evidence on the effects of disclosure comes predominantly from controlled experimental scenarios where the salience of the manipulation is artificially heightened [8,13], limiting the generalizability of their conclusions to the real advertising context.

Against this backdrop, the present study aims to address this knowledge gap through a field experiment evaluating the effect of spokesperson type (human versus AI-generated avatar) and authorship disclosure (stating or not stating the origin of the advertisement) on the performance metrics of an actual advertising campaign on the Meta platform. A 2×2 factorial design (spokesperson type $\times$ authorship disclosure) was implemented, stratified by gender and age group of the target audience, for a total of 384 observations. Performance metrics (reach, video plays, clicks, and messages) were modeled using a generalized structural equation model (GSEM) with a negative binomial distribution. This approach enables overcoming the limitations of self-report-based studies by capturing the actual behavior of the audience in response to advertising materials in their natural consumption environment.

The results reveal a paradoxical finding: the AI avatar succeeds in capturing the audiovisual attention of the audience and increasing video plays, but this advantage does not translate into greater interaction. On the contrary, reach—a delivery metric determined by the platform's algorithm rather than a deliberate consumer action—and, more markedly, clicks, which do involve active interaction, are significantly reduced compared to the human spokesperson, with spokesperson type exerting no effect whatsoever on messages.

This trade-off between attention and action constitutes a relevant empirical contribution to the debate on the effectiveness of synthetic spokespersons. Equally significant, authorship disclosure did not show a significant effect on any of the evaluated metrics, suggesting that transparency regarding the use of AI can be implemented without a detectable cost to campaign performance, although for messages the evidence remains inconclusive.

2. Theoretical Framework

2.1. Artificial Intelligence and AI Avatars in Digital Marketing

The capacity of AI to sustain predictive analysis of consumer behavior and the automated generation of synthetic content enables brands to anticipate market needs and optimize the return on advertising investment [14,15]. In this context, virtual influencers (VIs) emerge, defined as computer-generated characters with distinct personalities and narratives, designed to influence consumer behavior through digital platforms [3,6]. These avatars represent the culmination of diverse technological capabilities by integrating personalization, automation, and content generation into a single strategic asset [6]. The virtual influencer market is experiencing exponential growth driven by strategic advantages such as total control over the message, mitigation of reputational risks associated with human figures, uninterrupted availability for global campaigns, and greater long-term profitability [16,17]. Cases such as Lil Miquela and Lu do Magalu illustrate the capacity of these entities to generate brand recognition on an international scale [18]. Companies investing in these assets seek not only immediate advertising returns but also the construction of intellectual property aligned with the intangibles economy [4].

However, the effectiveness of avatars is not absolute and is mediated by consumer psychology. Despite advances in visual realism, evidence suggests that human spokespersons continue to be perceived as more credible and trustworthy [2,7,8]. Direct human interaction establishes a level of trust that technology has yet to fully replicate [19]. This relationship becomes more complex when considering that avatar type influences perceived credibility in a nonlinear fashion: hyperrealistic avatars may generate rejection due to the uncanny valley effect [20], whereas avatars with a fictitious appearance show no significant differences in credibility compared to human spokespersons [21]. This tension between the strategic advantages of avatars and consumer psychological barriers makes it essential to analyze their impact through theoretical frameworks that explain human perception of these artificial entities [3].

2.2. Effect of Spokesperson Type on Advertising Campaign Performance

Online advertising performance refers to the measurement of campaign effectiveness operationalized through a set of key performance indicators or KPIs [22]. On platforms such as Meta, these KPIs span the entire marketing funnel, from awareness metrics such as reach and impressions, visual engagement metrics such as video plays, to deliberate action metrics such as clicks and message conversions. This sequential hierarchy—exposure, visual attention, deliberate action, and conversion—aligns with the hierarchical model of advertising effects [23] and constitutes the organizing framework for the dependent variables of the present study.

Recent empirical evidence consistently shows that AI spokespersons and avatars generate greater visual attention than human spokespersons. AI involvement in advertisements evokes consumer curiosity, which increases engagement with advertising content [24], and eye-tracking evidence confirms that AI-generated advertisements capture visual attention comparable to or greater than advertisements with real models, particularly in fixation time and number of fixations [25]. A plausible explanation is that perceived novelty mediates the relationship between AI-generated advertising and behavioral intentions, and AI endorsers

with human-like appearance can evoke greater engagement than human endorsers and facilitate the acceptance of new products [26].

However, the visual attention captured by AI avatars does not necessarily translate into deliberate consumer actions. Evidence indicates that AI spokespersons reduce empathy and advertising effectiveness on action metrics, an effect attenuated only when the avatar exhibits elevated anthropomorphic cues [27]. The incongruence between virtual influencers and products requiring experiential evaluation equally reduces persuasive effectiveness [28]; for hedonic products, human influencers are significantly more persuasive than virtual ones [29]; and when the human-likeness and behavioral authenticity of virtual influencers are low, consumer behavioral responses are significantly reduced [30]. The convergence of these findings produces a paradoxical pattern: AI spokespersons capture more visual attention but generate fewer deliberate actions. AI labels appear to increase perceived novelty while simultaneously reducing authenticity, which is consistent with the coexistence of positive effects on attention metrics and negative effects on action metrics within the same advertising stimulus. Nevertheless, much of the existing evidence in this area tends to come from laboratory experiments and surveys relying on intention measures rather than actual behavior. Few studies have examined whether this pattern of greater visual attention and fewer deliberate actions replicates with real behavioral data from an advertising platform, documenting the simultaneous coexistence of both effects within the same field experiment. Based on the foregoing, the following hypotheses are proposed:

H1a. *Spokesperson type (AI avatar versus human) significantly influences the reach of the advertising campaign.*

H1b. *Spokesperson type (AI avatar versus human) significantly influences the video plays of the advertising campaign.*

H1c. *Spokesperson type (AI avatar versus human) significantly influences the clicks of the advertising campaign.*

H1d. *Spokesperson type (AI avatar versus human) significantly influences the messages generated by the advertising campaign.*

2.3. Effect of Authorship Disclosure on Advertising Performance

The decision to reveal that advertising content has been created or fronted by artificial intelligence introduces an additional layer of complexity that can alter campaign performance. Recent evidence suggests that AI authorship disclosure produces simultaneous and opposing effects: on one hand, it increases the perceived novelty of the content, which can improve attitudes toward the advertisement, and on the other, it may reduce perceived authenticity, thereby activating consumer skepticism.

The stream of research documenting negative disclosure effects is consistent. AI disclosure reduces advertising effectiveness, and brand strength can buffer this negative effect [31]. Convergently, disclosure of AI-generated content is generally understood to negatively affect consumer attitudes and behavioral intentions, with source type likely acting as a relevant moderating factor. Disclosure timing is also determinant, as early disclosure could reduce attitudes toward the advertisement when the appeal is rational, though this effect would plausibly be attenuated when the appeal is emotional.

Nevertheless, the disclosure effect is not universally negative. AI disclosures and media literacy determine the recognition of AI influencers, but this recognition does not always translate into significant behavioral changes, suggesting that disclosure may have limited effects when other visual cues already signal the artificial nature of the content. In

emerging markets, AI disclosure can significantly increase trust, which in turn positively mediates purchase intention [32].

The existing evidence, however, comes from attitudinal and intentional measures obtained in controlled experimental settings. There is little empirical evidence on the effect of AI authorship disclosure on real behavioral metrics within active advertising campaigns on social media. Moreover, existing studies evaluate disclosure in contexts where AI-generated content is not visually distinguishable from human content. Based on the foregoing, the following hypotheses are proposed:

H2a. *Explicit AI authorship disclosure will significantly affect the reach of the advertising campaign.*

H2b. *Explicit AI authorship disclosure will significantly affect the video plays of the advertising campaign.*

H2c. *Explicit AI authorship disclosure will significantly affect the clicks of the advertising campaign.*

H2d. *Explicit AI authorship disclosure will significantly affect the messages generated by the advertising campaign.*

Throughout H2 and H3, authorship disclosure denotes the presence of a statement of the advertisement's origin in the ad copy, as against its absence. Its wording is not identical across spokesperson conditions—"Contenido creado con IA" (content created with AI) in the avatar condition and "Elaboración propia" (own production) in the human condition—because in each case it states the actual origin of the stimulus; Section 3.2 details this operationalization and Section 6 discusses the limitation it entails.

2.4. Interaction Between Spokesperson Type and Authorship Disclosure

The interaction between spokesperson type and authorship disclosure constitutes a relevant aspect of the proposed model. The central prediction is that the disclosure effect is not uniform but operates differentially according to spokesperson condition. In the AI avatar condition, the artificial nature of the spokesperson is visually evident, as the avatar possesses stylistic characteristics that distinguish it from a real human. Therefore, verbal AI authorship disclosure provides information that the viewer can already infer visually. Evidence indicates that even subtle cues about the virtual nature of the influencer can reduce brand responses, although to a lesser extent than prominent disclosures [33], and that the modality of the AI cue—visual versus verbal—is a determining factor in disclosure impact, as disclosing that the image was AI-generated has different effects than disclosing that the text was AI-generated [34].

In contrast, in the human spokesperson condition the authorship statement declares a human origin (see Section 3.2), which the visual cues already support. The verbal cue is thus congruent with the stimulus, as it is in the avatar condition, yet it differs in informational status: it makes explicit an assumption the viewer would hold by default, rather than confirming a visible departure from it. This asymmetry in the relation between what is seen and what is declared could differentially modulate the effect on engagement metrics. Evidence suggests that the visual context modulates the effect of verbal disclosure and that strategic disclosure of AI involvement functions as a boundary condition whose impact varies according to consumer prior expectations [35]. Nevertheless, the cited studies have examined the interaction between disclosure format and content type in controlled experimental settings, but none has specifically tested whether verbal authorship disclosure differentially affects the human spokesperson and the avatar on real behavioral metrics

from active advertising campaigns. Based on the foregoing, the following hypotheses are proposed:

H3a. *Authorship disclosure will significantly moderate the effect of spokesperson type on the reach of the advertising campaign.*

H3b. *Authorship disclosure will significantly moderate the effect of spokesperson type on video plays.*

H3c. *Authorship disclosure will significantly moderate the effect of spokesperson type on the clicks of the advertising campaign.*

H3d. *Authorship disclosure will significantly moderate the effect of spokesperson type on the messages generated by the advertising campaign.*

2.5. Sociodemographic Variables: Receiver Gender and Age

Beyond the aforementioned variables that articulate the central hypotheses of the model, the literature on virtual influencer marketing identifies receiver gender and age as factors with independent explanatory and control capacity.

Regarding gender, evidence indicates that women and men differentially process credibility and authenticity signals from virtual influencers. The effects of visual and behavioral realism of the spokesperson on trust and purchase intention differ according to the VI's gender [35], and gender similarity between the receiver and the virtual influencer has direct effects on attitudes and behavioral intentions, with more intense responses when dimensions of homophily align simultaneously [36]. In the domain of parasocial relationships, the influence of these bonds on purchase intention is significantly stronger in women than in men [37]. However, meta-analytic evidence documents that women exhibit greater aversion to AI [34], suggesting that the effect of gender on advertising performance with avatars could be modulated by the interaction between greater emotional receptivity and greater technological resistance.

Regarding age, this variable articulates two influence mechanisms in potentially opposing directions. On one hand, younger generations exhibit greater familiarity with and acceptance of VIs [38,39], which could translate into higher engagement metrics under the AI avatar condition, reducing the effectiveness gap with the human spokesperson. On the other hand, Processing Speed Theory [40] posits that older consumers require more time to integrate new information, which could amplify or attenuate disclosure effects depending on the age group. In the context of legal education, the persuasive effect of digital avatars is more pronounced among older individuals, suggesting that format novelty may be more impactful in segments with less prior exposure [41]. Counter to intuition, millennials are more receptive to AI recommendations and disclosure labels than Generation Z, possibly because younger individuals are more desensitized to these technologies [42]. This complexity justifies the inclusion of both variables as controls with moderating potential in the experimental design.

It should be noted that the hypotheses of this study were formulated in a non-directional manner, that is, without anticipating the specific direction of the relationship between the independent variables and the advertising performance metrics. This methodological decision is grounded in three reasons. First, research on the impact of AI avatars on real behavioral metrics of advertising campaigns constitutes an emerging area where prior findings, predominantly based on self-report measures in controlled settings, are insufficient to predict with certainty the direction of effects in a real market environment. Second, the theoretical review reveals mixed or contradictory findings regarding the same phenomena across different contexts and platforms—avatars capture more visual attention

but generate fewer deliberate actions, and disclosure increases perceived novelty but reduces authenticity—which precludes confidently anticipating the direction of the effect. Third, the interaction between factors in tension—the particularities of Meta's delivery algorithm, the Latin American cultural context, and the divergent findings from other regions and platforms—generates an uncertainty that justifies not committing the research to directional predictions that could prove incorrect for at least one of the metrics analyzed.

The proposed research model (Figure 1) integrates these relationships by experimentally manipulating spokesperson type and authorship disclosure as independent variables, measuring their effects on four behavioral metrics—reach, video plays, clicks, and messages—and incorporating receiver gender and age group as control variables.

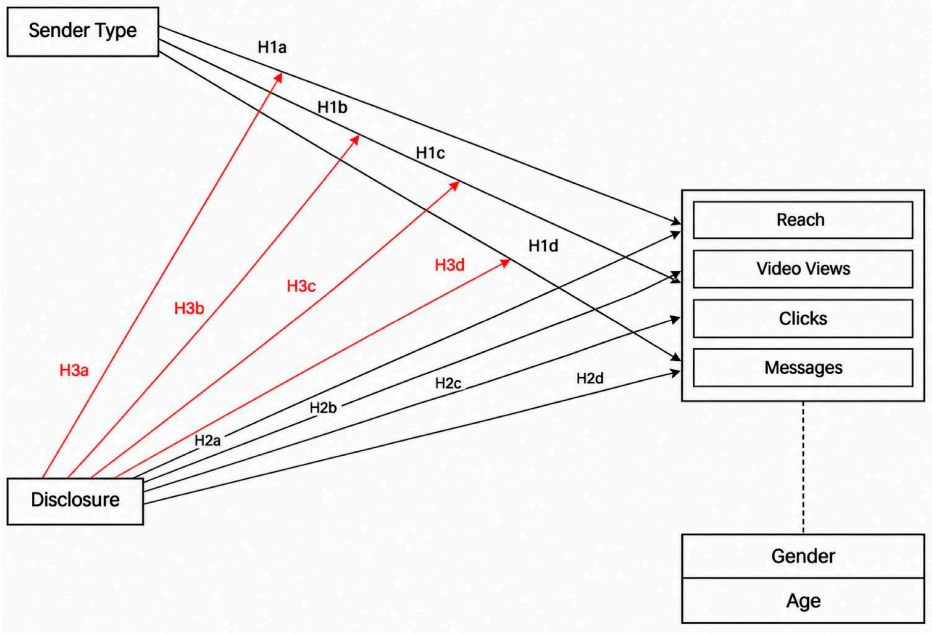


Figure 1. Research Model.

In summary, the hypotheses are organized into three groups following the logic of the advertising funnel: H1a–H1d predict differentiated effects of spokesperson type on each performance metric; H2a–H2d predict a significant effect of authorship disclosure; and H3a–H3d predict an interaction between both factors. Figure 1 integrates these relationships in the research model. The predictions are empirically evaluated in Section 4, where reminders of each hypothesis's content are incorporated alongside the corresponding results to facilitate reading.

3. Methodology

3.1. Approach and Experimental Design

To address the objectives of this research, an experimental approach was adopted, given that this method enables controlling extraneous variables and focusing specifically on the impact of spokesperson type on digital advertising performance metrics [43]. This precision and control enable a clearer interpretation of the direct effects of the spokesperson, a crucial aspect for understanding the influence of artificial intelligence (AI) on digital communication.

Unlike laboratory experiments based on self-report, this study implemented a field experiment through a paid advertising campaign on the Meta platform, enabling the capture of the audience's actual behavior in response to the advertising materials. A 2 × 2 factorial design was employed (spokesperson type: human spokesperson versus AI-generated

avatar; authorship disclosure: stating or not stating the origin of the advertisement in the ad copy), stratified by two demographic audience segmentation variables: gender (male or female) and age group (18 to 24 years or 25 to 60 years). The combination of the two experimental factors with the two stratification variables produced 16 cells, and each cell was replicated 24 times, for a total of 384 observations. Spokesperson type and authorship disclosure constitute the independent variables of interest whose interaction is explicitly modeled in the main GSEM. Gender and age group, although implemented through Meta Ads Manager's audience segmentation, are incorporated as main-effect covariates in the main model; their possible interactions with spokesperson type are examined as sensitivity analyses (see Section 3.5).

3.2. Stimuli and Procedure

The advertising context corresponded to a campaign for a university in southwestern Colombia. The message presented in the advertisements was as follows (see Box 1):

Box 1. Advertising Copy Used.

> Have you seen the new programs from Universidad del Valle EJECAFETERO? 👆 If you're in the Eje Cafetero or Valle del Cauca region, this is your chance to study a career with a real future in the region. 👇 Write to us and we'll tell you everything about admissions.

To isolate the effect of spokesperson type, stimuli were constructed in a paired manner through a four-stage generation process, each supported by a different AI tool. First, the human spokesperson was recorded delivering the original campaign message; from this recording, NotebookLM Standard was used to organize and refine the source content and generate a narrative summary (audio overview) that served as input for the script, without being directly integrated into the final advertisement. Second, using this input, Gemini Pro was employed to draft the final version of the script, design the visual image and appearance of the avatar, and generate the synthetic voice (text-to-speech) to replace the human spokesperson's voice. Third, Veo 3 was used to generate the video and avatar animation from said image and script, including lip synchronization (lip-sync) with the synthetic voice and the recreation of scenes and audiovisual elements from the original video. Finally, Vids was used both to complete the generation of the avatar as a virtual presenter and for the final assembly and editing of the advertisement (cuts, synchronization adjustments, and export). Through this process, the AI avatar version retained exactly the same message, script, and audiovisual elements as the human spokesperson version (see Figures 2 and 3). Thus, the message, script, and audiovisual elements were held constant across conditions, so that the manipulated dimension was the nature (human or synthetic) of the spokesperson. Matching does not, however, eliminate the residual differences inherent to the generation technology—voice naturalness, facial movement fluidity, appearance, and production quality—which remain confounded with spokesperson type and are addressed as a limitation in Section 6. The operationalization of the authorship disclosure variable is detailed in the following paragraph.

Both videos shared the same technical production specifications to ensure comparability across conditions. The duration of each advertisement was 35 s, with a resolution of 1080 × 1080 pixels in square format (1:1 aspect ratio) optimized for the Meta feed, a frame rate of 30 fps, and H.264 encoding. The delivery style in both versions corresponded to a spokesperson addressing the camera in a medium shot, presenting the academic offering in a direct and persuasive manner. The script, narrative structure, and duration of each segment were identical in both versions. Regarding voice characteristics, the human version employed the natural voice of the recorded person, while the avatar version used

a synthetic voice generated by Gemini Pro from the same script; speech rate, tone, and cadence were calibrated to replicate the prosodic characteristics of the human version. The avatar's lip synchronization was automatically generated by Veo 3 from the synthetic audio track. Nevertheless, despite efforts to match both versions, residual differences inherent to the generation technology—such as voice naturalness, facial movement fluency, or possible visual artifacts of the avatar—cannot be entirely ruled out and constitute a design limitation discussed in Section 6.



Figure 2. Human Spokesperson.



Figure 3. AI Avatar. Source: created using artificial intelligence.

It is necessary to specify the operationalization of the authorship disclosure variable. Disclosure was manipulated as the presence or absence of an authorship statement in the ad copy, the text accompanying the publication on the Meta platform and the content of that statement corresponded truthfully to the origin of each stimulus. In the human spokesperson disclosure condition, the copy included the phrase "Elaboración propia" (own production), indicating that the advertisement had been produced directly by the advertiser; in the AI avatar disclosure condition, the copy included the phrase "Contenido creado con IA" (content created with AI), indicating that both the spokesperson and the audiovisual content had been generated using artificial intelligence tools. In the non-disclosure conditions, no authorship statement of any kind appeared in the copy. This placement in the copy is consistent with disclosure practices for sponsored content on social media, where information about the origin of the content is communicated in the publication text [44,45]. The manipulation therefore contrasts the presence of a truthful provenance statement against its absence within each spokesperson condition; at no point was an advertisement accompanied by an authorship claim that did not correspond to its actual production process. Supplementary Material S2 reproduces the four advertisements as displayed on the platform, showing the complete copy and the position of the authorship

statement in the disclosure conditions; Figures 2 and 3 present representative frames of the human spokesperson and of the AI avatar, respectively. This operationalization underlies the interaction hypothesis (H3). Although both statements are truthful, they differ in their relation to the visual cues available to the viewer. In the avatar condition, the statement makes explicit an artificial origin that the stylistic characteristics of the avatar already suggest, so the verbal cue is largely redundant with the visual one. In the human spokesperson condition, the statement confirms an origin that the viewer would plausibly assume by default. It is therefore hypothesized that the effect of disclosure will differ significantly according to spokesperson condition.

Prior to the campaign launch, a pilot test was conducted with an independent sample of 20 university students from a program different from the one promoted in the campaign. Participants viewed the disclosure versions of both spokesperson conditions and were assessed on whether they identified the authorship statement included in the ad copy: 18 of 20 participants (90%) correctly identified it, indicating that the disclosure manipulation was perceptible and functioned as intended.

3.3. Advertising Campaign and Data Collection

The advertisements were disseminated on the Facebook platform, using performance as the campaign objective. Geographic segmentation was restricted to municipalities within the Coffee Cultural Landscape of Colombia (PCC) located in Valle del Cauca and Quindío, a territory that, despite being distributed across different departments, shares a marked cultural homogeneity centered on the coffee-growing tradition, paisa identity, and a common social and symbolic heritage [46]. This territorial delimitation enabled controlling for the effect of possible sociocultural differences between regions on performance metrics, thus isolating the variation attributable to experimental conditions. Segmentation by gender and age group was configured according to the experimental conditions.

The campaign was configured following the three-level hierarchical architecture of Meta Ads Manager: campaign → ad set → ad. A single campaign was created with a performance objective (conversions–messages), grouping all ad sets under a unified budget and optimization configuration. Within this campaign, 16 ad sets were configured, one for each scenario of the 2 × 2 factorial design, stratified by gender and age group. Each ad set defined a unique combination of demographic segmentation (gender × age group) and creative piece (spokesperson type × authorship disclosure). In this manner, gender and age segmentation in Meta generates mutually exclusive audiences at the ad set level: a user classified as a male aged 18–24 cannot simultaneously belong to the female aged 25–60 segment, which eliminates audience overlap between experimental cells defined by demographic factors [47].

The bidding strategy employed was lowest cost, which is Meta's default configuration and delegates bid optimization to the platform's algorithm to maximize results within the budget allocated to each advertisement. Placements were configured in automatic mode (Advantage+ placements), allowing the platform to distribute impressions across available Facebook placements (Feed, Stories, Reels, In-stream, Search) based on the estimated probability of conversion. This decision is consistent with methodological recommendations for field experiments on Meta platforms, which suggest maintaining standard algorithmic configuration to preserve the ecological validity of the experiment [47,48]. The formal A/B testing procedure of Meta (Experiments tool) was not employed, given that this tool only allows comparing a limited number of cells and does not support the 2 × 2 factorial design required by the study.

Randomization was implemented at two levels. At the first level, the assignment of the creative piece (spokesperson type × authorship disclosure) to each demographic combi-

nation (gender × age) was experimentally controlled: the 16 scenarios were defined a priori and each was replicated 24 times throughout the campaign. At the second level, within each ad set, the Meta platform randomly assigned eligible users to available impressions through its real-time bidding system, which selects which user to show each advertisement to based on an algorithm that combines the advertiser's bid, the estimated action rate, and ad quality. This mechanism has been documented as a valid form of quasi-experimental randomization in the literature on field experiments on digital platforms [49]. The scope of this randomization requires precision. What was randomly assigned were impressions to eligible users within each ad set, not users to creative conditions; the assignment of creative pieces to demographic cells was experimentally controlled and balanced, not randomized at the user level. Because demographic targeting produces mutually exclusive audiences across gender × age cells but not across the four creative conditions within a given cell, a user in the same demographic segment could in principle receive advertisements from more than one creative condition. Since delivery also depends on the bid, the estimated action rate, and ad quality, it cannot be ruled out that different creatives reached subaudiences with different response propensities. The estimates reported here are accordingly interpreted as performance differences under the platform's standard algorithmic delivery, rather than as causal effects arising from full random assignment of users to conditions.

Regarding audience overlap and observation independence, it is necessary to distinguish two types of possible competition. First, competition between conditions within the same demographic segment: within a given gender × age cell, the four creative conditions (human without disclosure, human with disclosure, avatar without disclosure, avatar with disclosure) were scheduled as independent advertisements with separate budgets. Given that each advertisement operated with a budget of COP 10,000 (USD 2.72) and reached an average of 4679 users per advertisement (see Section 4.1), while the eligible population in the study's geographic area—the PCC municipalities in Valle del Cauca and Quindío—amounts to approximately 726,692 inhabitants [50], the fraction of the audience exposed per advertisement was approximately 0.64% of the available population base (4678.80 ÷ 726,692). It should be noted that this proportion, while low, is not by itself sufficient to rule out cross-advertisement exposure, and that the effectively eligible audience is narrower than the territory's total population, since it excludes those who are not active platform users or who do not meet the targeting criteria. The stable unit treatment value assumption (SUTVA) required for causal inference [51] therefore rests on the within-advertisement exposure frequency (median = 1.04; minimum = 1.00) and on the delivery window of approximately 24 h, rather than on the population share; the possibility that a given user was exposed to more than one condition over the course of the campaign cannot be entirely ruled out and is acknowledged as a limitation in Section 6. Second, auction competition: although advertisements from the same segment participated in the same Meta auction pool, the budget per advertisement was sufficiently low that the competitive pressure between experimental conditions was minimal, a scenario analogous to that documented by Coppock et al. [47] in their field experiment with digital advertising on Facebook.

Regarding frequency and repeated exposure, the unit budget of COP 10,000 de facto restricted the duration of each advertisement to approximately 24 h before exhausting its budget, which limits the exposure window and the probability that a user views the same advertisement more than once. Additionally, Meta implements, by default, a frequency management system that diversifies delivery among unique users when the campaign objective is performance, minimizing individual overexposure. The median frequency observed in the campaign was 1.04 (minimum = 1.00; maximum = 1.11), confirming that the vast majority of impressions corresponded to single exposures.

The algorithmic learning phase of Meta, which typically requires approximately 50 optimization events per ad set, is completed rapidly with low budgets and performance objectives. Nevertheless, given that each individual advertisement constituted an independent observation with its own budget and delivery period, variations arising from the learning phase are randomly distributed across experimental conditions and do not introduce systematic bias. This design is methodologically equivalent to the "micro-experiments" approach with minimal budgets per cell documented in the literature on field experiments on social networks [52,53].

Each advertisement was placed with a fixed budget of COP 10,000 (approximately USD 2.72, at an average exchange rate of COP 3680 per USD during the campaign period [54]), so that each budget injection constituted an independent advertisement and, therefore, a study observation; in total, 24 advertisements were placed per scenario to complete the 384 observations. The total campaign investment amounted to COP 3,840,000 (approximately USD 1044) during the period from 14 May to 6 June 2026, distributed across 24 launch days. Sixteen advertisements were launched on each day, one per scenario, so that daily outlay was constant at COP 160,000. Spending was therefore identical by construction across the four experimental conditions: each condition comprised 96 advertisements and received COP 960,000. Read by factor, that same total splits into COP 1,920,000 per level of spokesperson type and, again, COP 1,920,000 per level of authorship disclosure. The observed differences therefore do not stem from unequal budget allocation. Placements were left on automatic (Advantage+), so the placement mix was determined by the platform; the campaign export does not disaggregate impressions by placement, so the comparability of the placement mix across conditions cannot be verified, a limitation acknowledged in Section 6. This order of magnitude is consistent with the methodological literature on low-cost field experiments on Meta platforms, which documents budgets of a few dollars per experimental cell as sufficient to generate controlled exposure and valid behavioral data, even in emerging economies comparable to Colombia [52,55]. For each advertisement, the Meta platform reported the performance metrics constituting the dependent variables of the study: reach, video plays, clicks, and messages. The study received approval from the ethics committee of Universidad del Valle.

To assess the possible influence of temporal effects on the results, the balance of the 16 experimental conditions was verified across the 24 launch days of the campaign. The design proved exactly balanced: each of the 16 scenarios appears exactly once on every launch day, so that the temporal structure is orthogonal by construction to the experimental conditions. The deposited dataset does not record the calendar date of each launch; it identifies the position of each launch day within its launch batch, a variable with 11 levels that group between one and three launch days each instead, so that their cell sizes differ. This is immaterial to the control they provide: because every launch day carried the 16 scenarios exactly once, any grouping of launch days is balanced with respect to the experimental conditions by construction, and each of the 11 levels accordingly contains the 16 scenarios in exactly the same proportion. The fixed effects reported below correspond to this variable, which is why their joint tests carry 10 degrees of freedom rather than 23. As a robustness analysis, a model with fixed effects for that variable was estimated. The day effects were jointly significant in three of the four equations (reach: $\chi^2(10) = 88.70$; video plays: $\chi^2(10) = 51.84$; clicks: $\chi^2(10) = 163.71$; all $p < 0.001$; messages: $\chi^2(10) = 16.32$, $p = 0.091$), reflecting day-to-day variation in the platform's delivery volume. As expected from a balanced design, however, their inclusion did not alter the experimental coefficients: reach b = −0.135 versus −0.138; video plays b = +0.224 versus +0.223; clicks b = −0.472 versus −0.476; messages b = −0.053 versus −0.053, and neither authorship disclosure nor the interaction reached significance in any equation. As a further check, robust standard

errors clustered by launch day (11 clusters) and by ad set (16 clusters) were estimated. The coefficients are identical by construction and the substantive conclusions hold under both: the three spokesperson effects on reach, video plays and clicks retain $p < 0.001$, and no disclosure or interaction effect reaches significance (p between 0.171 and 0.848 under day clustering and between 0.197 and 0.944 under ad-set clustering). Two discrepancies deserve mention. The interaction on reach moves from $p = 0.158$ with OIM standard errors and $p = 0.125$ with day fixed effects to $p = 0.002$ under day clustering, but returns to $p = 0.197$ under ad-set clustering; and the spokesperson effect on messages moves from $p = 0.513$ with OIM standard errors to $p = 0.097$ under ad-set clustering, without reaching significance. Because the cluster-robust variance estimator is known to be downward-biased with as few as 11 or 16 clusters and its asymptotic properties do not apply, and because the experimental factors do not vary within an ad set, neither discrepancy is interpreted as evidence of moderation or of a spokesperson effect on messages; day fixed effects are therefore reported as the primary temporal control and the two clustered specifications only as sensitivity checks [47]. Appendix A Table A10 reports the four specifications side by side for all experimental coefficients.

3.4. Operationalization of Variables

Table 1 summarizes the coding of the independent, control, and dependent variables used in the analysis.

Table 1. Variable Coding.

Variable	Type	Coding
Spokesperson type	Independent (manipulated)	0 = Human; 1 = AI Avatar
Authorship disclosure	Independent (manipulated)	0 = No disclosure; 1 = Discloses
Gender	Control	0 = Male; 1 = Female
Age group	Control	0 = 18–24 years; 1 = 25–60 years
Reach	Dependent	Count (non-negative integer)
Video plays (25%)	Dependent	Count (non-negative integer)
Clicks	Dependent	Count (non-negative integer)
Messages	Dependent	Count (non-negative integer)

It is necessary to specify the operational definitions of each metric as reported by the Meta Ads Manager API, given that the platform offers multiple variants of the same family of indicators. Reach corresponds to the number of unique individuals (distinct users) who saw the advertisement at least once; this metric is distinguished from impressions, which count the total number of times the advertisement was displayed on screen, including repeated views by the same user. Video plays at 25% records the total number of times—not unique—that the video was played for at least 25% of its duration (8.75 s for a 35-s video); this metric includes repeated plays and is activated regardless of whether the user initiated playback manually or whether the video played automatically in the feed (autoplay). Clicks correspond to Meta's "clicks (all)" metric, which counts all click interactions on the advertisement, including link clicks, reactions (like, love, etc.), comments, shares, clicks on the image or video, clicks to expand the copy text, and clicks on the advertiser's profile name. This distinction is relevant because the observed average of 555 clicks (all) per advertisement from 4869 impressions (rate of 11.4%) is consistent with typical rates of clicks (all) on Meta video advertisements but would be atypically high if it corresponded exclusively to link clicks, whose rates typically range between 0.5% and 2%. Finally, messages (messaging conversations started) record the number of messaging conversations initiated through the advertisement's call-to-action (CTA) button configured with the "Send message" objective; in the context of this higher education promotion campaign,

each message represents a lead or prospect who initiated direct contact with the institution, which is why it is employed as a conversion metric.

Regarding delivery metrics and reporting configuration, the attribution window employed was Meta's default configuration of 7 days after click and 1 day after view (7-day click, 1-day view), meaning that a conversion (message) is attributed to the advertisement if the user sent the message within 7 days of having clicked on the advertisement or within 24 h of having viewed it without clicking. The average frequency—mean number of times each unique individual saw the advertisement—was 1.04 (SD = 0.02), confirming that the vast majority of users were exposed only once to each advertisement. The average cost per thousand impressions (CPM) was COP 2054 (approximately USD 0.56), consistent with industry benchmarks for video campaigns in Latin American markets. All metrics were extracted directly from the Meta Ads Manager panel via CSV download on the day following the completion of each advertisement, without applying filters or additional adjustments to the data reported by the platform.

3.5. Empirical Model Specification

Given that the four dependent variables—reach, video plays at 25%, clicks, and messages—correspond to non-negative integer counts, models from the Poisson–negative binomial family were considered. Although the Poisson distribution is frequently used for count data, it assumes equidispersion (variance equal to the mean), an assumption that is frequently violated in practice, leading to underestimated standard errors and inflated Type I error rates [56]. The negative binomial distribution relaxes this assumption through an auxiliary dispersion parameter α that captures extra-Poisson variance [57]. To evaluate the distributional choice, likelihood ratio (LR) tests were conducted comparing each negative binomial equation with its Poisson restriction (H_0: $\alpha = 0$) (see Appendix A, Table A2).

The four outcome equations were integrated within a generalized structural equation model (GSEM). GSEM combines the features of generalized linear models and structural equation models, enabling the simultaneous treatment of continuous, binary, ordinal, or count variables within a unified estimation framework [58]. The choice of GSEM over four separate negative binomial regressions is grounded in three considerations: (i) joint maximum likelihood estimation produces a single variance-covariance matrix of parameter estimates, enabling cross-equation hypothesis testing without the need for post-hoc adjustments; (ii) a unified model provides a single log-likelihood, AIC, and BIC for overall fit evaluation and model comparison (see Appendix A, Table A4); and (iii) the framework is extensible to incorporate shared latent variables or additional structural paths if the data require them.

In the current specification, each outcome equation includes the same set of predictors—spokesperson type, authorship disclosure, their interaction, audience gender, and audience age group—estimated as independent negative binomial equations that do not share cross-equation error covariance or a latent factor. This independence assumption is substantively appropriate because the four metrics represent distinct stages of the advertising funnel (awareness, engagement, action, conversion) with different generating processes.

Formally, the model for the m-th outcome ($m = 1, \ldots, 4$) is specified as follows. The link function is the natural logarithm (log link), canonical for the negative binomial family:

$$\begin{aligned} ln(\mathrm{E}[y_{\mathrm{m,i}}|\mathrm{X_i}]) \quad &= \beta_{\mathrm{m},0} + \beta_{\mathrm{m},1}\cdot Spokesperson_{\mathrm{i}} + \beta_{\mathrm{m},2}\cdot disclosure_{\mathrm{i}} \\ &+ \beta_{\mathrm{m},3}\cdot(Spokesperson_{\mathrm{i}} \times disclosure_{\mathrm{i}}) + \beta_{\mathrm{m},4}\cdot Gender_{\mathrm{i}} \\ &+ \beta_{\mathrm{m},5}\cdot Age_{\mathrm{i}} \end{aligned} \tag{1}$$

where $y_{m,i}$ denotes the count for the m-th metric of the i-th advertisement, $\mu_{m,i} = E[y_{m,i} \mid X_i]$, and $\beta_{m,0}$ through $\beta_{m,5}$ are the regression coefficients specific to each equation. The distributional assumption is:

$$y_{m,i} \sim \text{BinNeg}(\mu_{m,i}, \alpha_m) \quad (2)$$

with variance function:

$$\text{Var}(y_{m,i}) = \mu_{m,i} + \alpha_m \cdot \mu_{m,i}^2 \quad (3)$$

where $\alpha_m > 0$ is the dispersion parameter specific to each equation. When $\alpha_m \to 0$, the negative binomial converges to the Poisson distribution. The coefficients are interpretable on the log-rate scale; exponentiation yields the incidence rate ratio (IRR):

$$\text{IRR}_k = \exp(\beta_{m,k}) \quad (4)$$

such that IRR_k represents the multiplicative change in the expected count associated with a one-unit change in X_k, holding the remaining predictors constant.

An exposure variable (offset) was not incorporated, as all advertisements received an identical fixed budget of COP 10,000, making raw counts directly comparable across observations. The Spokesperson × Disclosure interaction term was included to evaluate the potential moderating effect of authorship disclosure on spokesperson type (hypotheses H3a–H3d).

Nevertheless, given that Meta's delivery algorithm determines how many impressions and unique users each advertisement receives even with a fixed budget, it is necessary to distinguish between performance per unit of budget and the audience response conditioned on exposure. To this end, a complementary GSEM was estimated incorporating the natural logarithm of impressions as an exposure variable (offset) in the equations for video plays, clicks, and messages. In this offset model, the coefficients are interpreted as effects on rates—video view rate (VTR = plays/impressions), click-through rate (CTR = clicks/impressions), and message rate (messages/impressions)—rather than on raw counts. The reach equation was maintained without an offset, given that reach is itself a delivery metric rather than a conditioned response metric. Additionally, a descriptive advertising funnel analysis was conducted, disaggregated by experimental condition, decomposing the chain: budget → impressions → reach → video plays → clicks → messages.

The model was estimated by maximum likelihood (ML) via Newton-Raphson optimization, with standard errors derived from the observed information matrix (OIM). All analyses were conducted using Stata/SE 19.0 (StataCorp LLC, College Station, TX, USA). The complete replication code, anonymized campaign-level data, and audiovisual material used in the study are available in the data availability statement repository.

Given that only the messages variable exhibited zero counts (10 of 384 observations; 2.6%), zero-inflated negative binomial (ZINB) and hurdle specifications were estimated for this outcome and compared with the standard negative binomial via AIC and BIC (see Appendix A, Table A8). Finally, given that the study evaluates 12 primary hypotheses (H1a–d, H2a–d, H3a–d) across four outcomes, a Benjamini–Hochberg procedure controlling the false discovery rate at 5% was applied, along with a Bonferroni correction (adjusted $\alpha = 0.05/12 = 0.00417$), to assess the robustness of the significance pattern under multiple testing (see Appendix A, Table A7).

As an additional sensitivity analysis to evaluate whether gender and age group moderate the effects of spokesperson type, an extended GSEM specification was estimated that included the spokesperson × gender and spokesperson × age interactions in all four equations simultaneously. None of the eight interactions reached statistical significance after Benjamini–

Hochberg correction (all adjusted $p > 0.10$), indicating that the effects of spokesperson type on all four metrics do not differ significantly between males and females or between the 18–24 and 25–60 age groups. This result justifies the more parsimonious specification of the main model, which includes gender and age only as main-effect covariates.

4. Results

4.1. Descriptive Results

The descriptive statistics for quantitative variables are summarized in Table 2. On average, the advertisements registered 4869 impressions (SD = 962) and a reach of 4679 unique individuals (SD = 932). Regarding the behavioral dependent variables, an average of 58 plays at 25% (SD = 18.5), 555 clicks (SD = 164.5), and 3.9 messages or conversions per advertisement (SD = 2.2) were observed.

Table 2. Descriptive Statistics for Quantitative Variables.

Variable	Mean	SD	Minimum	Maximum	Median
Impressions	4868.88	962.03	1856	7436	4785.5
Reach	4678.80	931.55	1856	7238	4588.0
Video plays (25%)	57.94	18.47	9	120	55.0
Clicks	554.80	164.52	180	981	526.0
Messages	3.92	2.22	0	13	4.0

Table A1 presents the descriptive statistics disaggregated by the 16 cells of the design (2 spokesperson × 2 disclosure × 2 gender × 2 age), including the number of observations, mean, and standard deviation of impressions, reach, frequency, CPM, video plays at 25%, clicks (all), and messages, as well as the exposure-conditioned rates (VTR, CTR, and message-per-click rate) with their respective 95% confidence intervals. Figure 4 graphically represents the means and confidence intervals of the four dependent variables disaggregated by spokesperson type and disclosure condition, enabling visual assessment of the interaction magnitude and the uncertainty associated with each estimate.

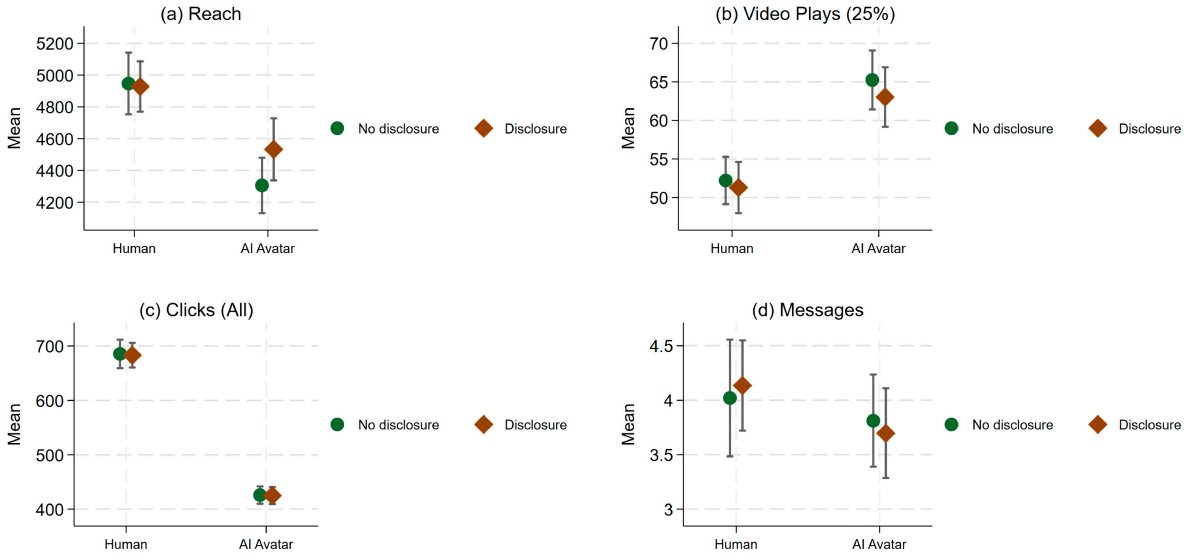


Figure 4. Means and 95% Confidence Intervals of the Four Dependent Variables by Spokesperson Type and Authorship Disclosure Condition.

Table 3 presents the distribution of categorical variables. Given the balanced factorial design, each level of the two experimental factors and two demographic stratification variables represents 50% of the observations.

Table 3. Descriptive Statistics for Qualitative Variables.

Variable	Category	Frequency	%
Spokesperson	Human	192	50.0
Spokesperson	AI Avatar	192	50.0
Authorship disclosure	No disclosure	192	50.0
Authorship disclosure	Disclosure	192	50.0
Gender	Male	192	50.0
Gender	Female	192	50.0
Age group	18–24 years	192	50.0
Age group	25–60 years	192	50.0

4.2. Empirical Analysis

Table 4 shows the results of the GSEM application with a negative binomial distribution. The dispersion parameter (ln alpha) was significant across all four equations, evidencing corrected overdispersion and suggesting that the negative binomial provides a better fit for the dataset than the Poisson distribution. Alongside each coefficient, the standard error (SE) and the incidence rate ratio (IRR) with its 95% confidence interval, obtained as the exponential of the coefficient, are reported to facilitate interpretation.

Table 4. GSEM Negative Binomial Empirical Model Results.

Equation/Variable	Coef. (b)	SE	IRR	95% CI (IRR)	*p*-Value
Reach (Dependent variable)					
Spokesperson: AI Avatar	−0.138	0.027	0.871	[0.826; 0.919]	0.000 ***
Disclosure: Discloses	−0.002	0.027	0.998	[0.946; 1.053]	0.938
Spokesperson x Disclosure	0.055	0.039	1.056	[0.979; 1.140]	0.158
Gender: Female	−0.049	0.020	0.952	[0.916; 0.990]	0.012 *
Age: 25–60	0.045	0.019	1.046	[1.007; 1.087]	0.019 *
Constant	8.506	-	-	-	0.000 ***
ln(alpha)	−3.330	-	a = 0.036	[0.031; 0.041]	***
Video plays (Dependent variables)					
Spokesperson: AI Avatar	0.223	0.043	1.250	[1.149; 1.361]	0.000 ***
Disclosure: Discloses	−0.018	0.044	0.982	[0.902; 1.070]	0.682
Spokesperson x Disclosure	−0.016	0.061	0.984	[0.873; 1.109]	0.789
Gender: Female	0.050	0.030	1.051	[0.990; 1.116]	0.101
Age: 25–60	−0.025	0.030	0.975	[0.919; 1.035]	0.409
Constant	3.942	-	-	-	0.000 ***
ln(alpha)	−2.634	-	a = 0.072	[0.060; 0.086]	***
Clicks (Dependent variables)					
Spokesperson: AI Avatar	−0.476	0.027	0.621	[0.589; 0.655]	0.000 ***
Disclosure: Discloses	−0.004	0.027	0.996	[0.945; 1.050]	0.893
Spokesperson x Disclosure	0.002	0.038	1.002	[0.930; 1.079]	0.965
Gender: Female	0.001	0.019	1.001	[0.964; 1.040]	0.954
Age: 25–60	0.012	0.019	1.012	[0.975; 1.051]	0.515
Constant	6.523	-	-	-	0.000 ***
ln(alpha)	−3.411	-	a = 0.033	[0.028; 0.038]	***

Table 4. *Cont.*

Equation/Variable	Coef. (b)	SE	IRR	95% CI (IRR)	*p*-Value
Messages (Dependent variables)					
Spokesperson: AI Avatar	−0.053	0.081	0.949	[0.810; 1.111]	0.513
Disclosure: Discloses	0.029	0.079	1.029	[0.881; 1.202]	0.716
Spokesperson x Disclosure	−0.059	0.114	0.942	[0.754; 1.178]	0.602
Gender: Female	0.090	0.057	1.095	[0.980; 1.223]	0.112
Age: 25–60	0.017	0.057	1.017	[0.910; 1.137]	0.767
Constant	1.336	-	-	-	0.000 ***
ln(alpha)	−2.892	-	a = 0.055	[0.025; 0.122]	***

Note: N = 384. IRR = exp(b). 95% CI calculated as exp(b ± 1.96 × SE). Base categories: Human, No disclosure, Male, 18–24 years. Standard errors from the observed information matrix (OIM). Estimation: ML via Newton-Raphson (Stata/SE 19.0, gsem command). Model fit: LL = −7911.56; AIC = 15,879.12; BIC = 15,989.73. * $p < 0.05$, *** $p < 0.001$. Of the 12 primary hypothesis tests (H1a–d, H2a–d, H3a–d), only the three significant spokesperson effects—on clicks, video plays, and reach—survive both the Bonferroni correction (adjusted $\alpha = 0.00417$) and the Benjamini–Hochberg procedure at FDR = 5%. The remaining nine, comprising all disclosure and interaction effects as well as the spokesperson effect on messages, do not reach significance under either adjustment (see Appendix A, Table A7).

The overall model yielded a log-likelihood of −7911.56 (AIC = 15,879.12; BIC = 15,989.73). Likelihood ratio (LR) tests comparing each negative binomial equation with its Poisson restriction overwhelmingly rejected the equidispersion assumption for all four outcomes (Reach: LR $\chi^2(1) = 60{,}934.32$, $p < 0.001$; Video plays: LR $\chi^2(1) = 944.50$, $p < 0.001$; Clicks: LR $\chi^2(1) = 5537.00$, $p < 0.001$; Messages: LR $\chi^2(1) = 8.38$, $p = 0.004$), confirming that the negative binomial provides a significantly better fit than the Poisson distribution (Appendix A, Table A2). The estimated dispersion parameters were: Reach $\alpha = 0.036$ [95% CI: 0.031; 0.041]; Video plays $\alpha = 0.072$ [0.060; 0.086]; Clicks $\alpha = 0.033$ [0.028; 0.038]; and Messages $\alpha = 0.055$ [0.025; 0.122]. The Pearson χ^2/df ratios for all four equations were close to 1.0 (range: 0.933–1.027), indicating adequate model fit. An extended model incorporating the 16 additional two-way interactions between experimental factors and sociodemographic controls did not improve fit (joint Wald $\chi^2 = 12.06$, df = 16, $p = 0.740$; ΔAIC = +20.03; ΔBIC = +83.25), supporting the parsimony of the reported specification (Appendix A, Tables A3 and A4).

4.2.1. Reach Path

In the reach path, it was found that the use of an AI avatar significantly reduces the number of unique individuals reached compared to the human spokesperson (b = −0.138; IRR = 0.87; $p < 0.001$), equivalent to an approximately 13% decrease. Authorship disclosure did not show a significant effect ($p = 0.938$), nor did its interaction with spokesperson type ($p = 0.158$). Among the control variables, the female audience was associated with lower reach (b = −0.049; $p = 0.012$), and the older age group with slightly higher reach (b = 0.045; $p = 0.019$). These results support H1a and lead to the rejection of H2a and H3a.

4.2.2. Video Plays Path

The video plays path revealed that the AI avatar significantly increases video plays compared to the human spokesperson (b = 0.223; IRR = 1.25; $p < 0.001$), representing an approximately 25% increase. Neither authorship disclosure ($p = 0.682$) nor the spokesperson × disclosure interaction ($p = 0.789$) was significant. The control variables also did not reach statistical significance (gender, $p = 0.101$; age, $p = 0.409$). H1b is supported and H2b and H3b are rejected.

4.2.3. Clicks Path

In the clicks path, it was observed that the AI avatar markedly and significantly reduces the number of clicks compared to the human spokesperson (b = −0.476; IRR = 0.62; $p < 0.001$), a decrease of approximately 38%. Neither authorship disclosure ($p = 0.893$) nor the spokesperson × disclosure interaction ($p = 0.965$) had a significant effect, nor did the control variables (gender, $p = 0.954$; age, $p = 0.515$). H1c is supported and H2c and H3c are rejected.

4.2.4. Messages Path

Finally, in the messages path, no significant effects were found for spokesperson type (b = −0.053; $p = 0.513$), authorship disclosure ($p = 0.716$), or their interaction ($p = 0.602$). The control variables were also not significant (gender, $p = 0.112$; age, $p = 0.767$). Consequently, H1d, H2d, and H3d are rejected.

Taken together, spokesperson type constitutes the only factor with a robust effect on performance, although its direction varies depending on the metric considered: the AI avatar increases video plays (+25%) but reduces reach (−13%) and, notably, clicks (−38%), with no effect on messages. This pattern of opposing signs reveals a trade-off: the AI avatar captures audiovisual attention but does not translate into clicks or greater reach, and even strongly penalizes click action. For its part, authorship disclosure did not show a statistically significant effect on any of the evaluated metrics. Equivalence testing further indicates that any disclosure effect is smaller than 10% for reach, video plays, and clicks, although for messages equivalence at that margin could not be established and the evidence remains inconclusive (Section 5; Appendix A, Table A6).

It is important to note that the raw metrics combine exposure opportunity (determined by Meta's delivery algorithm) with the conditioned audience response. To isolate the latter, a complementary model with an impressions offset was estimated (Section 4.3). The offset model results show that the differences in VTR and CTR persist after controlling for exposure volume, indicating that the trade-off between visual attention and deliberate action is not explained by the number of impressions received alone. Because the offset adjusts for exposure volume but not for audience composition, placement mix, or delivery timing, this analysis narrows—without fully eliminating—the role of the algorithmic delivery system.

As a sensitivity check for the messages equation—the only outcome exhibiting zero counts (10 of 384 observations; 2.6%)—zero-inflated negative binomial (ZINB) and hurdle models were estimated and compared with the standard negative binomial. Neither alternative improved model fit (NB: AIC = 1662.35, BIC = 1690.00; ZINB: AIC = 1671.52, BIC = 1722.88), confirming that the low proportion of zeros does not warrant a zero-inflated or hurdle specification (Appendix A, Table A8). Finally, to address the multiple testing problem arising from 12 primary hypothesis tests, both the Bonferroni correction (adjusted $\alpha = 0.05/12 = 0.00417$) and the Benjamini–Hochberg procedure (FDR = 5%) were applied. Under both corrections, three of the twelve tests—the spokesperson effects on reach, video plays, and clicks (H1a, H1b, H1c; all $p < 0.001$)—remained significant, while the remaining nine retained their non-significance (Appendix A, Table A7). The study's pattern of significance is therefore robust to adjustment for multiple comparisons, although it should be made explicit that three, not twelve, hypotheses survive the adjustment.

4.3. Analysis of Exposure-Conditioned Rates and Advertising Funnel

To address the possibility that the observed effects on raw counts reflect differences in exposure opportunity—determined by Meta's delivery algorithm—rather than differences in consumer response, a complementary GSEM was estimated incorporating the natural

logarithm of impressions as an offset (exposure variable) in the equations for video plays, clicks, and messages. In this model, the coefficients are interpreted as effects on response rates per impression.

The results of the offset model confirm the conclusions of the main model: the AI avatar significantly increases the video view rate (VTR; IRR = 1.434, $p < 0.001$) and significantly reduces the click-through rate (CTR; IRR = 0.707, $p < 0.001$), while the message rate does not differ significantly between conditions (IRR = 1.080, $p = 0.318$). Authorship disclosure and the spokesperson × disclosure interaction were not significant in any of the offset equations, replicating the findings of the non-offset model. These results indicate that the observed differences in raw counts are not explained by the volume of impressions received alone. The offset nevertheless adjusts for the amount of exposure and not for the composition of the audience reached, the placement mix, or delivery timing—dimensions that the platform's optimization may also introduce and that the campaign export does not allow to be disaggregated; the analysis therefore narrows the role of algorithmic delivery without eliminating it (see Appendix A, Table A9).

Table 5 presents a descriptive advertising funnel analysis disaggregated by spokesperson type and authorship disclosure. Each funnel metric explicitly identifies its numerator, denominator, and aggregation method: mean impressions per advertisement under a fixed budget of COP 10,000; Reach/Impr. = reach ÷ impressions × 100; VTR = video plays at 25% ÷ impressions × 100; CTR = clicks (all) ÷ impressions × 100; and Msg/Click = messages ÷ clicks (all) × 100. All rates are computed per advertisement and then averaged across the advertisements of each condition, so the reported values are means of rates rather than ratios of means. These definitions are identical to those of the deposited replication file and of Table A1, Panel B. This analysis enables distinguishing at which stage of the funnel the experimental conditions diverge.

Table 5. Advertising Funnel Analysis by Experimental Condition.

Condition	n	Impressions	Reach/Impr. (%)	VTR (%)	CTR (%)	Msg/Click (%)
Human—No disclosure	96	5136.32 (997.11)	96.38 (2.03)	1.03 (0.29)	13.45 (1.77)	0.58 (0.33)
Human—Disclosure	96	5125.96 (828.79)	96.21 (1.87)	1.00 (0.29)	13.40 (1.54)	0.61 (0.29)
AI Avatar—No disclosure	96	4508.44 (892.09)	95.51 (2.18)	1.48 (0.46)	9.50 (0.87)	0.90 (0.49)
AI Avatar—Disclosure	96	4704.81 (978.12)	96.28 (1.98)	1.35 (0.34)	9.11 (0.91)	0.88 (0.50)
Total	384	4868.88 (962.03)	96.09 (2.04)	1.22 (0.40)	11.36 (2.46)	0.74 (0.44)
p-value (spokesperson)	—	<0.001	0.004	<0.001	<0.001	<0.001

Note: Cell entries are means with standard deviations in parentheses; $n = 96$ advertisements per condition (384 total). Fixed budget of COP 10,000 per advertisement. Definitions: Reach/Impr. = reach ÷ impressions × 100, the share of impressions delivered to distinct individuals (its reciprocal is the exposure frequency, mean 1.04). VTR = video plays at 25% ÷ impressions × 100; the numerator counts total plays rather than unique viewers and includes autoplay and repeated plays, so this rate indexes exposure to video content rather than deliberate viewing. CTR = clicks (all) ÷ impressions × 100; the numerator is Meta's "clicks (all)" metric, which aggregates link clicks with reactions, comments, shares, and image, video, and profile-name clicks, and is therefore an order of magnitude larger than a link-click rate. Msg/Click = messages ÷ clicks (all) × 100. All rates are computed per advertisement and then averaged within each condition. *p*-values test the spokesperson main effect in a Gaussian GLM with robust (sandwich) standard errors controlling for disclosure, the spokesperson × disclosure interaction, gender, and age. Rates with 95% confidence intervals disaggregated by all 16 design cells appear in Appendix A, Table A1, Panel B. The cell means, the Total row and these *p*-values are reproduced in Supplementary Material S1 on the deposited dataset.

The funnel analysis reveals that the differences between the human spokesperson and the AI avatar manifest differentially along the conversion chain. At the delivery stage, advertisements with the AI avatar received approximately 10% fewer impressions per budget unit than those with the human spokesperson (4607 versus 5131 mean impressions), while the share of impressions delivered to distinct individuals differed by less than one

percentage point across conditions (95.5–96.4%), a difference that nonetheless reaches statistical significance (p = 0.004; Table 5). This difference in delivery efficiency is what motivates the impressions-offset model. At the exposure-conditioned stages, the AI avatar shows a higher VTR (1.35–1.48% versus 1.00–1.03%) and a lower CTR (9.11–9.50% versus 13.40–13.45%) than the human spokesperson, a pattern that holds in the offset model. These results indicate that the observed differences are not explained by exposure volume alone; because the offset does not adjust for audience composition, placement mix, or delivery timing, the analysis narrows the role of algorithmic delivery without ruling it out. It should further be noted that VTR is built on total plays—which include autoplay and repeat plays—and that CTR is built on "clicks (all)", so both rates index exposure and broad interaction rather than deliberate viewing or movement toward enrollment. Authorship disclosure does not generate significant differences at any stage of the funnel.

5. Discussion

The results of the GSEM with a negative binomial distribution demonstrate that spokesperson type constitutes the only factor with a robust and significant effect on advertising performance, although its direction varies depending on the metric considered. This finding provides partial support for the first group of hypotheses (H1a–H1d) and reveals a trade-off between the AI avatar's capacity to capture audiovisual attention and its inability to translate that attention into behavioral actions of greater depth in the conversion funnel. For its part, authorship disclosure exerted no significant effect on the evaluated metrics, leading to the rejection of hypotheses H2a–H2d and H3a–H3d.

The positive effect of the AI avatar on video plays, which supports H1b, is consistent with evidence on the capacity of virtual influencers to attract greater initial visual attention [1]. The scope of this metric should be made explicit: video plays at 25% count total plays rather than unique users, including autoplay in the feed as well as repeat plays, so they constitute an indicator of exposure to video content rather than direct evidence of deliberate viewing, curiosity, or sustained attention. With that caveat, the observed increase is consistent with the idea that the avatar operates as a novel stimulus that interrupts the automatic processing of advertising content in the feed and raises the probability that the video is played. The Uncanny Valley Theory [20] offers a complementary explanation: the avatar's appearance, sufficiently human to be intriguing but not enough to generate full rejection, could position itself in a zone of perceptual curiosity that stimulates visual exploration, a mechanism consistent with the notion that human-likeness constitutes the central antecedent of interaction with virtual influencers [59], although in this case the interaction materializes in the viewing phase and does not extend to subsequent funnel stages.

By contrast, the AI avatar significantly reduces both reach and clicks, supporting H1a and H1c. The drop in clicks is particularly pronounced and suggests that the visual curiosity generated by the avatar does not translate into the trust necessary for the consumer to execute a deliberate action such as clicking on the advertisement. This pattern is consistent with Source Credibility Theory [60]: while the avatar can project a high level of visual attractiveness, its fabricated nature compromises the trustworthiness dimension that is critical for activating behaviors requiring greater cognitive commitment [6]. The processing fluency associated with the perception of the human spokesperson as a natural and familiar entity facilitates more favorable evaluations and stronger behavioral intentions [1], which would explain the advantage of the human spokesperson on metrics requiring active consumer decision-making. In the specific context of higher education promotion, where institutional trust and spokesperson expertise are critical conditions for activating contact intention [34,61], the avatar's disadvantage in clicks is consistent with the experiential nature of the promoted service.

It should be noted that, given that reach is a delivery metric determined largely by Meta's auction and distribution algorithm, the observed reduction could reflect not only lower audience receptivity but also differences in cost per thousand impressions (CPM) or in the efficiency of algorithmic distribution for advertisements featuring the avatar. However, the complementary analysis with an impressions offset (Section 4.3) shows that, even after controlling for exposure volume, the AI avatar's CTR remains significantly lower than that of the human spokesperson, indicating that the drop in clicks is not explained by the volume of impressions received alone; because the offset does not adjust for audience composition, placement mix, or delivery timing, this result narrows the role of differential delivery without ruling it out. The distinction between performance per budget and conditioned response to exposure strengthens the interpretation that the avatar captures visual attention (higher VTR) but fails to mobilize deliberate consumer action (lower CTR), regardless of how many exposure opportunities it receives.

The absence of a significant effect of spokesperson type on messages, leading to the rejection of H1d, suggests that the decision to send a message represents a qualitatively distinct conversion level from the preceding metrics. Sending a message implies a deliberate and personalized contact intention that is likely mediated by factors such as genuine interest in the educational offering, perceived need, and the receiver's personal circumstances—variables that operate independently of spokesperson nature. This result converges with evidence indicating that the effectiveness of influencer type varies according to the depth of the metric evaluated and that conversions involving greater commitment depend more on product relevance than on spokesperson characteristics [62].

Hypotheses H2a–H2d, which predicted a significant effect of authorship disclosure on the four performance metrics, were not supported by the data: none of the coefficients associated with disclosure reached statistical significance. It is important to emphasize that the absence of statistical significance does not constitute evidence of the nonexistence of the effect. The 95% confidence intervals for the disclosure IRR include values both below and above 1.0 across all four equations, indicating that the data are compatible with small effects in both directions. Post-hoc power analysis indicates that with the available sample (N = 384) and an alpha level of 0.05, the study has 80% power to detect an IRR of approximately 1.15 or 0.87; smaller effects cannot be ruled out with the current data. To evaluate whether the data are consistent with practical equivalence between the disclosure and non-disclosure conditions, TOST (two one-sided tests) equivalence tests were conducted with two margins set a priori: $\pm$20% (IRR between 0.82 and 1.22), corresponding to the smallest effect of practical interest for advertising campaign management, and $\pm$10% (IRR between 0.90 and 1.11), used as a more demanding benchmark. At the $\pm$20% margin, equivalence was established across all four metrics. At the $\pm$10% margin, equivalence was established for reach (both one-sided $p < 0.001$), video plays ($p = 0.003$ and $p = 0.030$), and clicks (both $p < 0.001$), but not for messages, where the upper one-sided test did not meet the threshold ($p = 0.184$; lower $p = 0.052$) and the evidence remains inconclusive (Appendix A, Table A6). It should be emphasized that a failure to establish equivalence indicates insufficient precision to rule out an effect of that magnitude, not the presence of one. Nevertheless, this result must be interpreted with caution given that the authorship statement was presented in the advertisement copy in a specific cultural context (Colombia) and cannot be generalized without replications to other formats, placements, or markets.

The practical implication of this result is substantive: organizations can incorporate AI authorship disclosures in their advertising campaigns, given that the results do not provide evidence that disclosure deteriorates advertising performance within the detection limits of this design—with the caveat that message equivalence at $\pm$10% was not established and the result remains inconclusive—suggesting that compliance with emerging regulatory

frameworks on algorithmic transparency could be implemented without an apparent cost to commercial effectiveness. Algorithm aversion [63,64] and the machine heuristic [65], although theoretically plausible as penalization mechanisms, did not manifest in a real Meta advertising context, suggesting that these effects may be more associated with laboratory experimental contexts where manipulation salience is artificially heightened.

The rejection of interaction hypotheses H3a–H3d indicates that authorship disclosure does not moderate the effect of spokesperson type on any of the evaluated metrics. This result means that the moderation hypothesis found no empirical support, which is not equivalent to demonstrating that the mechanisms activated by spokesperson nature and by disclosure operate independently: the equivalence tests establish equivalence of the interaction at $\pm 10\%$ for clicks only, so for reach, video plays, and messages the evidence is inconclusive (Appendix A, Table A6). It should further be noted that, because the interaction term is included in the model, the disclosure coefficient in Table 4 estimates the effect of the authorship statement in the reference condition, that is, in the human spokesperson. The corresponding linear contrasts show that the effect of the authorship statement likewise fails to reach significance within the avatar condition (reach IRR = 1.054 [0.999; 1.112], $p = 0.055$; video plays IRR = 0.966 [0.889; 1.051], $p = 0.424$; clicks IRR = 0.998 [0.946; 1.053], $p = 0.943$; messages IRR = 0.970 [0.826; 1.139], $p = 0.708$), nor as an average effect across both spokesperson conditions (reach IRR = 1.026, $p = 0.192$; video plays IRR = 0.974, $p = 0.394$; clicks IRR = 0.997, $p = 0.884$; messages IRR = 0.999, $p = 0.987$). The absence of detectable moderation is theoretically relevant because it runs counter to the prediction that the combination of an AI avatar with explicit disclosure would constitute the most unfavorable condition, as the lower inherent credibility of the artificial spokesperson converges with the intensified activation of persuasion knowledge [10,66]. The most parsimonious explanation is that authorship disclosure, by failing to effectively activate the PKM in this context, lacks the capacity to interact with the spokesperson variable to produce a differential effect.

The patterns observed in this study—greater visual attention but less deliberate action with the AI avatar, and absence of a significant disclosure effect—are consistent with several theoretical interpretations that should be noted as speculative given that the corresponding mediating constructs were not measured. The attention–action trade-off could be explained by stimulus novelty [67] insofar as the avatar constitutes a visually novel stimulus that captures attention but does not generate the connection necessary for deliberate action. Similarly, Source Credibility Theory [60] would suggest that avatars are perceived as less credible, reducing deliberate engagement, and the Uncanny Valley Theory [20] would propose that quasi-human avatars generate discomfort that inhibits action. The absence of a significant disclosure effect could be interpreted through the Persuasion Knowledge Model [68] or the machine heuristic, but these causal pathways cannot be confirmed with the available data. Future research should incorporate the measurement of these mediators to disentangle the specific psychological pathways underlying the behavioral patterns documented in this investigation.

In the main model (Table 4), gender and age group were significant only for reach. However, these differences must be interpreted with caution because reach is a delivery metric determined largely by Meta's auction algorithm. The lower reach observed for the female audience and for the 25-to-60-year age group could reflect differences in auction prices (CPM), in available advertising inventory, or in competition for those segments within the platform, rather than psychological differences in advertising information processing. Given that the budget was fixed per observation, a higher CPM for certain demographic segments mechanically translates into a lower number of impressions and, consequently, lower reach, without necessarily involving a differential behavioral mechanism. The offset model, in which each metric is expressed per impression, is consistent with this reading

rather than contrary to it. There the female audience, which receives fewer impressions, responds at a higher rate (VTR IRR = 1.100, $p = 0.001$; CTR IRR = 1.051, $p < 0.001$; message rate IRR = 1.149, $p = 0.011$), while the 25-to-60 group, which receives more impressions, shows slightly lower video play and click rates (IRR = 0.929, $p = 0.007$; IRR = 0.967, $p = 0.002$) (Appendix A, Table A9). The demographic differences therefore arise on the exposure side of the funnel and do not indicate lower responsiveness of these segments once exposure is held constant.

6. Conclusions

The findings of this study demonstrate that advertising spokesperson type constitutes a factor associated with digital campaign performance, although its effect is not unidirectional. The AI avatar is associated with a significant increase in video plays (+25%) and with a significant reduction in reach (−13%) and clicks (−38%). This pattern—greater visual attention but less deliberate action—is tentatively interpreted as an attention–action trade-off, which persists in the impressions offset model (Section 4.3), indicating that the observed differences are not explained by exposure volume alone; because the offset does not adjust for audience composition, placement mix, or delivery timing, this result narrows the role of Meta's algorithmic distribution without ruling it out. For their part, the disclosure hypotheses were not supported: the results of the equivalence tests suggest that, within the limits of the present design, stating the origin of the advertisement did not produce a performance deterioration exceeding 10% for reach, video plays, and clicks; for messages, equivalence at ±10% was not established and the evidence remains inconclusive (Appendix A, Table A6).

These findings have practical implications in two domains. From the perspective of campaign management, marketing managers seeking to maximize message visibility and video content consumption could benefit from the strategic incorporation of AI avatars, particularly in upper-funnel campaigns aimed at generating brand awareness. Conversely, the human spokesperson showed a significant and large advantage in clicks (−38% for the avatar), which is the deliberate-interaction metric measured in this study. It should be noted, however, that this advantage did not extend to messages—the only conversion metric in the design—where the effect of spokesperson type was not statistically significant ($p = 0.513$), and power was low (0.10). The data therefore support an advantage of the human spokesperson in click interaction, but leave the question of conversion effectiveness unresolved, which would require designs with greater power on that metric. A hybrid strategy employing AI avatars in the initial contact phases and human spokespersons in the conversion phases could capitalize on the differentiated strengths of each spokesperson type. Regarding transparency, the fact that disclosure did not significantly deteriorate performance is encouraging for organizations considering the adoption of voluntary transparency practices, although it does not constitute sufficient evidence to draw direct inferences about compliance with regulatory frameworks such as the European Union AI Act, whose scope, disclosure format, and cultural context differ substantially from the experimental conditions of the present study.

Nevertheless, these results must be interpreted with caution in light of several limitations. First, the study was confined to a single platform (Meta/Facebook) and a specific advertising context—the promotion of higher education—which limits the generalizability of findings to other platforms and product categories. Meta's algorithmic distribution mechanisms could condition the results in ways that would not be replicable on platforms with different distribution architectures, such as TikTok or YouTube. Second, the campaign targeted a specific geographic region in southwestern Colombia, introducing a possible cultural specificity in audience responses that may not be representative of other markets.

Third, the observation window comprised 24 launch days, across which the 16 experimental conditions were exactly balanced; while this rules out temporal confounding of the experimental contrasts, it does not allow capturing medium- or long-term effects such as advertising fatigue or the evolution of attitudes toward AI avatars with repeated exposure. Fourth, the choice of a fixed budget of COP 10,000 per observation, while ensuring comparability across conditions, restricts the total reach of each individual advertisement and could limit the variance observed in conversion metrics, particularly in messages. Additionally, the pilot test was limited to verifying the perceptibility of the authorship statement with a small sample (n = 20, distributed across the two disclosure versions) and did not include a formal assessment of perceptual comparability between the stimuli in dimensions such as production quality, message clarity, or visual attractiveness; although both videos were produced with the same technical specifications, uncontrolled residual differences between the human and avatar versions could have influenced performance metrics beyond the intended experimental manipulation. Fifth, the design employed a single exemplar of the human spokesperson and a single exemplar of the AI avatar, which implies that spokesperson type is partially confounded with the idiosyncratic characteristics of the specific exemplars used. This stimulus sampling limitation implies that the findings are based on the specific creatives tested in this study and in the particular context of a higher education campaign, and therefore their generalization should be approached with caution. Relatedly, both spokespersons were unknown to the audience, which precludes evaluating the moderating role of celebrity, given that meta-analytic evidence confirms that the celebrity effect influences attitudes and purchase intention. Finally, the experimental design does not permit disentangling the intermediate psychological mechanisms (perceived credibility, processing fluency, persuasion knowledge activation) that mediate between the manipulated variables and performance metrics, given that exclusively aggregate behavioral indicators were employed. Four further clarifications about the design should be added to the limitations above. First, although the authorship statement was truthful in both spokesperson conditions, its wording necessarily differed between them, so the content of the statement is confounded with spokesperson type. The two formulations are not semantically symmetrical: the declaration of AI authorship names a specific production technology and is salient in current public debate, whereas the attribution of own production is a conventional credit formula that does not explicitly deny the use of AI and may be read as a generic authorship note rather than as a provenance disclosure. The main effect of disclosure (H2) therefore averages two manipulations that are unlikely to be equivalent in salience, and the interaction (H3) contrasts the effect of declaring AI authorship in the avatar condition against the effect of declaring own production in the human condition, rather than the effect of a single statement applied across formats. A design that crosses spokesperson type with three copy conditions—no statement, AI authorship statement, and human authorship statement—would separate the effect of stating provenance from the effect of what is stated and from the congruence between the two, and constitutes a natural extension of this work. Second, the pilot test verified that the authorship statement was identifiable under pilot conditions, but it does not establish equivalent visibility across the different placements through which the platform distributed the advertisements (Feed, Stories, Reels, In-stream, and search), nor does the campaign export allow impressions to be disaggregated by placement in order to verify the comparability of that mix across conditions. Third, random assignment operated on impressions within each ad set rather than on the assignment of users to creative conditions, so it cannot be entirely ruled out that a given user was exposed to more than one condition, or that the platform's optimization directed different creatives toward subaudiences with different response propensities; the estimated effects are accordingly interpreted as performance differences under standard

algorithmic delivery. Fourth, perceived novelty, credibility, and persuasion knowledge are invoked throughout the discussion as possible mechanisms, but they were not measured in this design and cannot be confirmed with the available behavioral indicators.

Future lines of research could address these limitations in several ways. The use of multiple exemplars of each spokesperson type would allow separating the effect of spokesperson type from the particular characteristics of each exemplar. The inclusion of spokespersons with different levels of fame would allow evaluating the moderating role of celebrity, given that recent evidence suggests that emotional responses differ substantially between human celebrities and AI virtual spokespersons [69]. Replication of the design across other platforms and product categories would allow evaluating the robustness of the observed effects and the possible moderation by product type—utilitarian versus experiential. The incorporation of process measures through complementary laboratory studies, such as perceived credibility, persuasion knowledge activation, and processing fluency, would enable the identification of the causal mechanisms underlying the attention–action trade-off documented in this research. Likewise, longitudinal designs evaluating the evolution of the AI avatar's effect on performance across multiple exposures would help determine whether the advantage in video plays is maintained or erodes with familiarity, and whether the disadvantage in clicks attenuates as consumers become accustomed to synthetic spokespersons. Research on the interaction between avatar type (hyperrealistic, animated, stylized) and product context would shed light on what degree of human-likeness is optimal for each campaign objective, contributing to a more sophisticated management of artificial intelligence-based communication assets. Finally, research with multiple disclosure formats, platforms, and markets is needed to evaluate the generalizability of the transparency findings.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/jtaer21100337/s1. S1, the reproducible Stata code used for the analysis, which is available through the following Zenodo repository: https://doi.org/10.5281/zenodo.22850471, supplementary_S1_postestimation_round2.do (Stata post-estimation script for the second revision round) and supplementary_S1_output.log (its complete output). Running S1 on the deposited dataset reproduces the linear contrasts reported in Section 5, the temporal balance and spending figures reported in Section 3.3, the funnel rates and spokesperson p-values of Table 5, and specifications (1) to (3) of Appendix A, Table A10; specification (4), with standard errors clustered by ad set, is reproduced by the deposited replication script. S2, supplementary_S2_advertisements.jpg, reproduces the four advertisements as displayed on the Meta platform, showing the full copy, the position of the authorship statement in the disclosure conditions, and the research-participation notice.

Author Contributions: Conceptualization, P.A.L.-H.; methodology, P.A.L.-H.; software, P.A.L.-H.; validation, E.A.-E. and A.R.O.; formal analysis, E.A.-E.; investigation, P.A.L.-H.; data curation, E.A.-E. and P.A.L.-H.; writing—original draft preparation, E.A.-E. and P.A.L.-H.; writing—review and editing, E.A.-E. and A.R.O.; visualization, A.R.O.; supervision, A.R.O.; project administration, E.A.-E. and A.R.O.; funding acquisition, E.A.-E. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Universidad del Valle, grant number CI 9211.

Institutional Review Board Statement: The study was conducted in accordance with the Declaration of Helsinki and approved by the Ethics Committee for Research of the Faculty of Administration Sciences (CEICA), Universidad del Valle (Acta No. 2, approved on 12 March 2026).

Informed Consent Statement: The advertisements informed users that they were part of a research study. No personally identifiable information was collected, and all analyses were based exclusively on aggregated, anonymized performance metrics provided by the Meta platform (e.g., reach, video plays, clicks, and messages).

Data Availability Statement: The original data presented in this study are openly available in Zenodo at https://doi.org/10.5281/zenodo.22699736 (reference number 22699736). The post-estimation script used in this revision round—linear contrasts of the disclosure and spokesperson effects, day-of-launch fixed effects, standard errors clustered by launch day and by ad set, the balance and spending checks reported in Section 3.3, and the funnel rates and *p*-values of Table 5—is provided as Supplementary Material S1, together with its Stata output log. It operates on the same deposited dataset and does not re-specify the model.

Acknowledgments: In the creation of this article, Anthropic's Claude 5 Opus generative artificial intelligence was used to improve the orthotypographic quality of the manuscript content in English.

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

Appendix A

Table A1. Descriptive statistics by factorial design cell.

Spokesperson	Disclosure	Gender	Age	n	Impressions	Reach	Video Plays (25%)	Clicks (All)	Messages	Freq.	CPM	VTR (%) [95% CI]	CTR (%) [95% CI]	Msg/Click (%) [95% CI]
Human	No disclosure	Male	18–24	24	5251.75 (1164.14)	5052.83 (1101.31)	49.79 (12.43)	685.83 (158.17)	3.50 (2.60)	1.04	2052.61	0.97 [0.87; 1.07]	13.02 [12.46; 13.59]	0.48 [0.35; 0.62]
Human	No disclosure	Male	25–60	24	5625.79 (991.54)	5460.00 (963.37)	49.67 (13.55)	685.08 (124.33)	3.79 (1.98)	1.03	1834.23	0.89 [0.80; 0.97]	12.21 [11.78; 12.63]	0.55 [0.44; 0.65]
Human	No disclosure	Female	18–24	24	4644.75 (747.06)	4461.42 (717.64)	53.54 (17.28)	681.38 (119.36)	4.46 (3.24)	1.04	2212.94	1.16 [1.01; 1.32]	14.76 [13.92; 15.60]	0.65 [0.47; 0.82]
Human	No disclosure	Female	25–60	24	5023.00 (825.05)	4815.21 (767.28)	55.79 (16.95)	689.58 (121.01)	4.33 (2.65)	1.04	2048.10	1.11 [1.00; 1.21]	13.79 [13.13; 14.46]	0.63 [0.49; 0.77]
Human	Discloses	Male	18–24	24	5202.92 (725.53)	4988.25 (683.04)	52.38 (15.39)	683.17 (108.01)	4.08 (2.28)	1.04	1957.22	1.00 [0.90; 1.11]	13.13 [12.62; 13.65]	0.60 [0.46; 0.74]
Human	Discloses	Male	25–60	24	5093.17 (878.91)	4895.71 (811.57)	46.25 (16.32)	676.38 (115.03)	4.00 (1.82)	1.04	2023.99	0.92 [0.79; 1.05]	13.39 [12.64; 14.13]	0.59 [0.48; 0.71]
Human	Discloses	Female	18–24	24	4791.42 (769.62)	4615.58 (751.91)	52.12 (17.68)	680.08 (109.74)	4.54 (1.98)	1.04	2144.41	1.09 [0.95; 1.24]	14.29 [13.58; 14.99]	0.66 [0.55; 0.77]
Human	Discloses	Female	25–60	24	5416.33 (856.66)	5213.96 (800.33)	54.38 (15.86)	692.33 (120.97)	3.92 (2.15)	1.04	1890.45	1.00 [0.90; 1.10]	12.79 [12.33; 13.25]	0.57 [0.44; 0.70]
AI Avatar	No disclosure	Male	18–24	24	4298.50 (873.73)	4094.38 (835.59)	65.58 (19.29)	420.33 (79.55)	3.75 (2.19)	1.05	2418.49	1.57 [1.35; 1.79]	9.83 [9.45; 10.20]	0.88 [0.69; 1.07]
AI Avatar	No disclosure	Male	25–60	24	4791.21 (987.74)	4586.42 (961.84)	63.08 (19.59)	433.54 (88.32)	3.62 (1.91)	1.05	2176.91	1.34 [1.17; 1.51]	9.08 [8.77; 9.40]	0.85 [0.65; 1.06]
AI Avatar	No disclosure	Female	18–24	24	4259.08 (734.49)	4079.54 (735.10)	68.79 (18.65)	418.71 (72.91)	4.00 (1.89)	1.05	2415.94	1.64 [1.44; 1.83]	9.85 [9.53; 10.16]	0.97 [0.76; 1.18]
AI Avatar	No disclosure	Female	25–60	24	4684.96 (884.18)	4465.17 (830.37)	63.54 (18.68)	430.62 (75.74)	3.88 (2.44)	1.05	2208.85	1.37 [1.21; 1.54]	9.24 [8.89; 9.60]	0.91 [0.68; 1.14]
AI Avatar	Discloses	Male	18–24	24	4945.08 (1088.84)	4777.46 (1058.87)	65.75 (22.15)	426.21 (76.81)	3.17 (2.01)	1.04	2121.55	1.33 [1.19; 1.46]	8.71 [8.37; 9.06]	0.74 [0.54; 0.93]

Table A1. *Cont.*

Spokesperson	Disclosure	Gender	Age	n	Impressions	Reach	Video Plays (25%)	Clicks (All)	Messages	Freq.	CPM	VTR (%) [95% CI]	CTR (%) [95% CI]	Msg/Click (%) [95% CI]
AI Avatar	Discloses	Male	25–60	24	4691.29 (976.42)	4501.62 (959.67)	60.21 (17.84)	423.58 (76.37)	4.00 (1.59)	1.04	2221.93	1.29 [1.16; 1.43]	9.10 [8.72; 9.49]	0.97 [0.80; 1.13]
AI Avatar	Discloses	Female	18–24	24	4690.38 (1002.81)	4516.67 (991.67)	61.54 (15.72)	419.00 (77.26)	3.58 (2.00)	1.04	2225.16	1.33 [1.20; 1.46]	9.00 [8.68; 9.32]	0.83 [0.67; 1.00]
AI Avatar	Discloses	Female	25–60	24	4492.50 (836.31)	4336.58 (843.91)	64.62 (20.53)	431.04 (82.52)	4.04 (2.46)	1.04	2301.87	1.44 [1.27; 1.61]	9.61 [9.21; 10.02]	0.97 [0.68; 1.25]

Note: n = 24 advertisements per cell (384 total). Columns up to CPM report means with standard deviations in parentheses; the three rate columns report means with 95% confidence intervals in brackets. Freq. = exposure frequency (impressions ÷ reach). CPM = cost per thousand impressions in Colombian pesos (COP 10,000 ÷ impressions × 1000). VTR = video view-through rate (25% views/impressions × 100). CTR = click-through rate (clicks/impressions × 100). Msg/Click = message-per-click rate (messages/clicks × 100). Fixed budget of COP 10,000 per advertisement.

Table A2. Justification of the negative binomial: LR comparison with Poisson, dispersion, and goodness of fit.

Dependent Variable	LR chi2(1) vs. Poisson	*p*	alpha [95% CI]	Pearson chi2/df
Reach	60,934.32	<0.001	0.036 [0.031; 0.041]	1.001
25% Views	944.50	<0.001	0.072 [0.060; 0.086]	1.012
Clicks	5537.00	<0.001	0.033 [0.028; 0.038]	0.933
Messages	8.38	0.004	0.055 [0.025; 0.122]	1.027

Note: $N = 384$. Likelihood ratio tests compare each negative binomial equation with its Poisson restriction ($\alpha = 0$). Rejection of equidispersion in all four outcomes supports the negative binomial specification. Pearson χ^2/df values near 1.0 indicate adequate fit of the negative binomial.

Table A3. Extended factorial model: Wald tests (2 × 2 × 2 × 2 design).

Contrast	chi2	df	*p*
Joint Wald: 16 interactions	12.06	16	0.740
Reach	8.40	4	0.078
25% Views	1.74	4	0.783
Clicks	0.31	4	0.989
Messages	1.60	4	0.808

Note: $N = 384$. The extended specification adds the spokesperson × gender, spokesperson × age, disclosure × gender, and disclosure × age interactions to all four equations simultaneously (16 additional parameters). None of the demographic moderation tests reaches significance, supporting the parsimonious specification reported in Table 4.

Table A4. Fit comparison: parsimonious vs. extended.

Model	LL	AIC	BIC
Original (28 parameters)	−7911.56	15,879.12	15,989.73
Extended (44 parameters)	−7905.58	15,899.15	16,072.98

Table A5. Post-hoc statistical power (alpha = 0.05, two-tailed).

Effect	Dep. Variable	b	SE	|z|	*p*	Power
AI Avatar	Reach	−0.1377	0.0274	5.03	<0.001	0.999
AI Avatar	25% Views	+0.2234	0.0431	5.19	<0.001	0.999
AI Avatar	Clicks	−0.4762	0.0270	17.66	<0.001	1.000
AI Avatar	Messages	−0.0527	0.0805	0.65	0.513	0.100
Disclosure	Reach	−0.0021	0.0274	0.08	0.938	0.051
Disclosure	25% Views	−0.0178	0.0436	0.41	0.682	0.069
Disclosure	Clicks	−0.0036	0.0268	0.14	0.893	0.052
Disclosure	Messages	+0.0288	0.0792	0.36	0.716	0.065
Interaction	Reach	+0.0548	0.0387	1.41	0.158	0.293
Interaction	25% Views	−0.0163	0.0610	0.27	0.789	0.058
Interaction	Clicks	+0.0017	0.0381	0.04	0.965	0.050
Interaction	Messages	−0.0594	0.1139	0.52	0.602	0.082

Table A6. TOST equivalence tests for non-significant effects (log scale).

Effect	Dep. Variable	p (Upper)	p (Lower)	Equiv.?	Conclusion
Limits +/−0.20 (IRR: 0.82–1.22)					
Disclosure	Reach	<0.001	<0.001	Yes	Equivalence established
Disclosure	25% Views	<0.001	<0.001	Yes	Equivalence established
Disclosure	Clicks	<0.001	<0.001	Yes	Equivalence established
Disclosure	Messages	0.015	0.002	Yes	Equivalence established
Interaction	Reach	<0.001	<0.001	Yes	Equivalence established
Interaction	25% Views	<0.001	0.001	Yes	Equivalence established
Interaction	Clicks	<0.001	<0.001	Yes	Equivalence established
Interaction	Messages	0.011	0.109	No	Inconclusive
Limits +/−0.10 (IRR: 0.90–1.11)					
Disclosure	Reach	<0.001	<0.001	Yes	Equivalence established
Disclosure	25% Views	0.003	0.030	Yes	Equivalence established
Disclosure	Clicks	<0.001	<0.001	Yes	Equivalence established
Disclosure	Messages	0.184	0.052	No	Inconclusive—not equivalent at ±10%
Interaction	Reach	0.122	<0.001	No	Inconclusive
Interaction	25% Views	0.028	0.085	No	Inconclusive
Interaction	Clicks	0.005	0.004	Yes	Equivalence established
Interaction	Messages	0.081	0.361	No	Inconclusive

Note: N = 384. Two one-sided tests (TOST) on the log-rate scale, applied to the coefficients of Table 4. Equivalence is declared only when both one-sided p-values fall below 0.05, which corresponds to the 90% confidence interval lying entirely inside the margin. The ±0.20 margin (IRR 0.82–1.22) reflects the smallest effect size of practical interest for campaign management, and the ±0.10 margin (IRR 0.90–1.11) provides a more demanding benchmark; both were set a priori. At the ±10% margin, equivalence is established for the disclosure effect on reach, video plays, and clicks, but not for its effect on messages, where the evidence remains inconclusive; the interaction is equivalent only for clicks. A failure to establish equivalence indicates insufficient precision to rule out an effect of that magnitude, not the presence of one.

Table A7. Multiple testing correction for the 12 hypotheses (Bonferroni and Benjamini–Hochberg).

Hypothesis	p (Unadjusted)	Bonferroni	B-H
AI Avatar -> clicks	<0.001	Yes	Yes
AI Avatar -> 25% views	<0.001	Yes	Yes
AI Avatar -> reach	<0.001	Yes	Yes
Interaction -> reach	0.158	No	No
AI Avatar -> messages	0.513	No	No
Interaction -> messages	0.602	No	No
Disclosure -> 25% views	0.682	No	No
Disclosure -> messages	0.716	No	No
Interaction -> 25% views	0.789	No	No

Table A7. *Cont.*

Hypothesis	p (Unadjusted)	Bonferroni	B-H
Disclosure -> clicks	0.893	No	No
Disclosure -> reach	0.938	No	No
Interaction -> clicks	0.965	No	No

Note: N = 384; m = 12 hypothesis tests (3 effects × 4 outcomes), ranked by ascending p-value. Bonferroni threshold: $\alpha = 0.05/12 = 0.00417$. Benjamini–Hochberg step-up critical value: (rank/12) × 0.05, controlling the false discovery rate at 5%. Three of the twelve hypotheses survive both corrections—the spokesperson effects on clicks, video plays, and reach. The remaining nine, including all disclosure and interaction effects and the spokesperson effect on messages, do not reach significance under either adjustment.

Table A8. Zero-inflation assessment for messages: simple NB vs. ZINB.

Model	AIC	BIC	Conclusion
Simple negative binomial	1662.35	1690.00	Selected specification
ZINB (zero inflation)	1671.52	1722.88	Does not improve fit

Table A9. GSEM with exposure offset: incidence rate ratios conditioned on impressions.

Predictor	25% Views IRR [95% CI]	p	Clicks IRR [95% CI]	p	Messages IRR [95% CI]	p	Reach IRR [95% CI]	p
AI Avatar	1.434 [1.329; 1.548]	<0.001	0.707 [0.687; 0.728]	<0.001	1.080 [0.928; 1.257]	0.318	0.871 [0.826; 0.919]	<0.001
Disclosure	0.977 [0.904; 1.055]	0.551	0.998 [0.969; 1.027]	0.868	1.029 [0.886; 1.194]	0.709	0.998 [0.946; 1.053]	0.938
Sender × Disclosure	0.938 [0.842; 1.045]	0.248	0.961 [0.922; 1.002]	0.063	0.903 [0.729; 1.119]	0.351	1.056 [0.979; 1.140]	0.158
Female	1.100 [1.042; 1.161]	0.001	1.051 [1.030; 1.074]	<0.001	1.149 [1.032; 1.279]	0.011	0.952 [0.917; 0.989]	0.012
25–60 years	0.929 [0.880; 0.980]	0.007	0.967 [0.947; 0.988]	0.002	0.973 [0.874; 1.083]	0.611	1.046 [1.007; 1.087]	0.019

Note: N = 384. Negative binomial equations with ln(impressions) as an exposure offset for video plays, clicks, and messages, so coefficients are interpreted as effects on rates per impression; the reach equation is reported without an offset, because reach is itself a delivery metric rather than a response conditioned on exposure. Base categories: Human, No disclosure, Male, 18–24 years. Conditioning on impressions removes differences in exposure volume between conditions, but it does not adjust for differences in audience composition, placement mix, or delivery timing that the platform's optimisation may also introduce; the offset results therefore narrow, without fully eliminating, the role of algorithmic delivery.

Table A10. Temporal robustness: day-of-launch fixed effects and clustered standard errors.

Outcome	Predictor	b (1) OIM	p (1) OIM	b (2) Day FE	p (2) Day FE	p (3) Day-Clustered	p (4) Ad-Set-Clustered
Reach	AI Avatar	−0.1377	<0.001	−0.1352	<0.001	<0.001	<0.001
	Disclosure	−0.0021	0.938	−0.0012	0.962	0.848	0.944
	Spokesperson × Disclosure	0.0548	0.158	0.0537	0.125	0.002	0.197
25% Views	AI Avatar	0.2234	<0.001	0.2240	<0.001	<0.001	<0.001
	Disclosure	−0.0178	0.682	−0.0176	0.669	0.171	0.536
	Spokesperson × Disclosure	−0.0163	0.789	−0.0164	0.775	0.473	0.668

Table A10. *Cont.*

Outcome	Predictor	b (1) OIM	*p* (1) OIM	b (2) Day FE	*p* (2) Day FE	*p* (3) Day-Clustered	*p* (4) Ad-Set-Clustered
Clicks	AI Avatar	−0.4762	<0.001	−0.4724	<0.001	<0.001	<0.001
	Disclosure	−0.0036	0.893	−0.0030	0.895	0.532	0.496
	Spokesperson × Disclosure	0.0017	0.965	0.0005	0.987	0.873	0.841
Messages	AI Avatar	−0.0527	0.513	−0.0532	0.501	0.485	0.097
	Disclosure	0.0288	0.716	0.0278	0.720	0.708	0.512
	Spokesperson × Disclosure	−0.0594	0.602	−0.0584	0.601	0.535	0.351

Note: N = 384. Log-scale coefficients of the GSEM negative binomial system reported in Table 4. Specification (1) is the reported model with OIM standard errors; specification (2) adds fixed effects for the position of the launch day within its batch (11 levels, each grouping between one and three launch days and identical in scenario composition), which is the temporal variable contained in the deposited dataset; specifications (3) and (4) keep the reported model and replace only the variance estimator with standard errors clustered by launch day (11 clusters) and by ad set (16 clusters), respectively, so that their coefficients are identical to (1) by construction. Joint Wald tests of the day effects in (2): reach $\chi^2(10) = 88.70$; video plays $\chi^2(10) = 51.84$; clicks $\chi^2(10) = 163.71$ (all $p < 0.001$); messages $\chi^2(10) = 16.32$ ($p = 0.091$). Because the cluster-robust estimator is downward-biased with as few as 11 or 16 clusters, and because the experimental factors do not vary within an ad set, specification (2) is the primary temporal control and (3) and (4) are reported only as sensitivity checks. Specifications (1) to (3) are reproduced by Supplementary Material S1 on the deposited dataset; specification (4) is reproduced by the deposited replication script, which estimates the same system with vce(cluster escenario).

References

1. Du, M.; You, K.H. Persuasive Differences between Human and Virtual Influencers in Health Supplement Advertising: Evidence from Eye-Tracking. *Front. Psychol.* **2026**, *16*, 1692737. [CrossRef] [PubMed]
2. Seymour, M.; Yuan, L.; Dennis, A.; Riemer, K. Facing the Artificial: Understanding Affinity, Trustworthiness, and Preference for More Realistic Digital Humans. In Proceedings of the 53rd Hawaii International Conference on System Sciences (HICSS), Maui, HI, USA, 7–10 January 2020.
3. Audrezet, A.; Koles, B.; Guidry Moulard, J.; Ameen, N.; McKenna, B. Virtual Influencers: Definition and Future Research Directions. *J. Bus. Res.* **2025**, *200*, 115647. [CrossRef]
4. Ahumada, H.; Chisari, O.; De Pablo, J.C.; Elías, V.J.; Fanelli, J.M.; Gasparini, L.; Heymann, D.; Montuschi, L.; Navajas, F. *Efectos de la Inteligencia Artificial (IA) en la Economía y el Análisis Económico*; Academia Nacional de Ciencias Económicas (ANCE): Buenos Aires, Argentina, 2021.
5. Saura, J.R.; Škare, V.; Ozretić Došen, D. Is AI-Based Digital Marketing Ethical? Assessing a New Data Privacy Paradox. *J. Innov. Knowl.* **2024**, *9*, 100597. [CrossRef]
6. Miao, F.; Kozlenkova, I.; Wang, H.; Xie, T.; Palmatier, R. An Emerging Theory of Avatar Marketing. *J. Mark.* **2022**, *86*, 67–90. [CrossRef]
7. Vallis, C.; Wilson, S.; Gozman, D.; Buchanan, J. Student Perceptions of AI-Generated Avatars in Teaching Business Ethics: We Might Not Be Impressed. *Postdigit. Sci. Educ.* **2023**, *6*, 537–555. [CrossRef]
8. Ozdemir, O.; Kolfal, B.; Messinger, P.R.; Rizvi, S. Human or Virtual: How Influencer Type Shapes Brand Attitudes. *Comput. Hum. Behav.* **2023**, *145*, 107771. [CrossRef]
9. Powers, G.; Johnson, J.P.; Killian, G. To Tell or Not to Tell: The Effects of Disclosing Deepfake Video on US and Indian Consumers' Purchase Intention. *J. Interact. Advert.* **2023**, *23*, 339–355. [CrossRef]
10. Wortel, C.; Vanwesenbeeck, I.; Tomas, F. Made with Artificial Intelligence: The Effect of Artificial Intelligence Disclosures in Instagram Advertisements on Consumer Attitudes. *Emerg. Media* **2024**, *2*, 547–570. [CrossRef]
11. Seeger, F.; Wessel, M.; Lehrer, C. AI Content Labeling and User Engagement in Social Media: The Role of AI Level, Content Type, and Disclosure Timing. *Electron. Mark.* **2026**, *36*, 31. [CrossRef]
12. Thomas, V.L.; Fowler, K. Close Encounters of the AI Kind: Use of AI Influencers as Brand Endorsers. *J. Advert.* **2021**, *50*, 11–25. [CrossRef]
13. Karaca, Ş.; Yildirim, E. Disclosure Matters: Perceived Manipulation, Perceived Ethics, and Purchase Intention toward AI Influencers in Social Media Marketing. *J. Theor. Appl. Electron. Commer. Res.* **2026**, *21*, 194. [CrossRef]

14. Islam, M.A.; Fakir, S.I.; Masud, S.B.; Hossen, M.D.; Islam, M.T.; Siddiky, M.R. Artificial Intelligence in Digital Marketing Automation: Enhancing Personalization, Predictive Analytics, and Ethical Integration. *Edelweiss Appl. Sci. Technol.* **2024**, *8*, 6498–6516. [CrossRef]
15. Al Khaldy, M.A.; Al-Obaydi, B.A.A.; Al Shari, A.J. The Impact of Predictive Analytics and AI on Digital Marketing Strategy and ROI. In *Cutting-Edge Business Technologies in the Big Data Era*; Yaseen, S.G., Ed.; Springer: Cham, Switzerland, 2023. [CrossRef]
16. Fast Company México. Los Avatares con IA Son el Futuro del Marketing, y Esta Compañía ya los Emplea en México. *Yahoo Noticias*, 28 May 2025.
17. Portada. El Marketing de Influencers y la Inteligencia Artificial. *Portada*, 28 May 2024.
18. Carrillo-Durán, M.V.; García García, M.; Corzo Cortés, L. Influencers Virtuales de Apariencia Humana como Forma de Comunicación Online: El Caso de Lil Miquela y Lu do Magalu en Instagram. *Rev. Comun.* **2024**, *23*, 3453. [CrossRef]
19. Gerlich, M. The Power of Virtual Influencers: Impact on Consumer Behaviour and Attitudes in the Age of AI. *Adm. Sci.* **2023**, *13*, 178. [CrossRef]
20. Mori, M.; MacDorman, K.F.; Kageki, N. The Uncanny Valley [from the Field]. *IEEE Robot. Autom. Mag.* **2012**, *19*, 98–100. [CrossRef]
21. Osorio-Andrade, C.F.; Arango-Espinal, E.; Arango-Pastrana, C.A. De lo Real a lo Ficticio: Evaluación de la Credibilidad de Noticias Difundidas por Humanos y por Avatares Creados con Inteligencia Artificial. *Palabra Clave* **2024**, *27*, e2738. [CrossRef]
22. Lou, C.; Yuan, S. Influencer Marketing: How Message Value and Credibility Affect Consumer Trust of Branded Content on Social Media. *J. Interact. Advert.* **2019**, *19*, 58–73. [CrossRef]
23. Park, S.; Chu, Y.; Chung, J.-H. Designing AI Agents for New Product Endorsement: Do Human-Like or Cartoon-Like AI-Generated Endorsers Evoke More Positive Ad Engagement from Consumers? *J. Curr. Issues Res. Advert.* **2025**, *46*, 259–277. [CrossRef]
24. van Berlo, Z.M.C.; Breves, P.L. Disclosing the Virtual Nature of Virtual Influencers: The Effect of Disclosure Prominence and the Role of Product Digitality. *Comput. Hum. Behav. Rep.* **2025**, *19*, 100742. [CrossRef]
25. Lavidge, R.J.; Steiner, G.A. A Model for Predictive Measurements of Advertising Effectiveness. *J. Mark.* **1961**, *25*, 59–62. [CrossRef]
26. Kiper, E.; Liang, Y. Do AI-Generated Fashion Ads Work as Effectively as Real Fashion Ads? An Eye-Tracking Comparison of Consumers' Visual Attention. *J. Glob. Fash. Mark.* **2026**, *17*, 1–17. [CrossRef]
27. Hou, W.; Xu, J.; Ren, J. Exploring the Advertising Effectiveness of Generative Artificial Intelligence: An Empirical Study Using the Hierarchy-of-Effects Model. *Int. J. Advert.* **2026**, *45*, 6–26. [CrossRef]
28. Yim, A.; Liska, L.; Mu, Y.; Thakkar, M. I Feel No Empathy toward AI: The Effects of AI vs Human Spokespersons on Advertisement Effectiveness. *J. Res. Interact. Mark.* **2026**, *20*, 108. [CrossRef]
29. Rejón-Guardia, F.; Palomas-Gómez, V.; Molinillo, S. Effects of Product Type on Persuasion Mechanisms of Human and Virtual Influencers. *Int. J. Consum. Stud.* **2026**, *50*, e70169. [CrossRef]
30. Bakr, Y.; Meshreki, H.; ElKashif, M. A Comparative Analysis of Virtual and Human Influencers in Sensory Endorsements: A Schema Congruity and Cue Diagnosticity Theory Perspective. *Technol. Soc.* **2026**, *90*, 103357. [CrossRef]
31. Shi, Y.; Jiang, Z. Consumer Responses to AI Disclosure Labels: The Role of Novelty and Authenticity. *SAGE Open* **2026**, *16*, 21582440261417793. [CrossRef]
32. Nguyen, B.P.; Wu, W. How and When Do Virtual Influencers Work? A Meta-Analysis of Mechanisms and Moderators in Digital Commerce. *J. Theor. Appl. Electron. Commer. Res.* **2026**, *21*, 124. [CrossRef]
33. Ma, J.; Zhang, D.; Chen, C.; Du, H.S. Matching Generative AI Word-of-Mouth with Product Type: Impact on Consumer Adoption and Trust. *J. Retail. Consum. Serv.* **2026**, *89*, 104615. [CrossRef]
34. Li, S.; Han, X.; Liu, H.; Zhang, M. "Virtual" Humans but "Real" Emotions? The Impact of Emotional Expression of Virtual Influencers on Users' Attitudes toward Tourist Destinations. *Tour. Manag. Perspect.* **2026**, *61*, 101457. [CrossRef]
35. Xiao, Z.; Wu, G. From Disclosure to Control: A Responsible AI Storytelling Model for CSR Promotions. *J. Promot. Manag.* **2026**, *32*, 240–251. [CrossRef]
36. Luo, L.; Wang, K.; Zheng, Y. Emojis as Catalysts: How AI Agents Leverage Emojis to Enhance Customer Engagement across the Pre-Core and Core Service Stages. *J. Retail. Consum. Serv.* **2026**, *89*, 104635. [CrossRef]
37. Zaharia, S.; Asici, J. Exploring the Impact of Virtual Influencers on Social Media User's Purchase Intention in Germany: An Empirical Study. *Lect. Notes Comput. Sci.* **2024**, *14720*, 108–126. [CrossRef]
38. Becaro-Lapiz, J.B.; Besario, M.R.A.; Bueno, N.L.E.; Bonghanoy, C.L.; Navaja, R.I.; Basilisco, G.L.; Sumilhig, J.M.; Ople-Alviola, C. AI Virtual Influencers in the Eyes of Future Professionals: Knowledge, Perceptions, and Risks among Filipino Students. In Proceedings of the 2025 11th International Conference on Education and Technology (ICET 2025), Chongqing, China, 26–28 September 2025; pp. 56–60. [CrossRef]
39. Hernández Serrano, M.J.; Renés Arellano, P.; Lena Acebo, F.J.; González Larrea, B. From Integration to Interference: Adolescents' Perceived Effects of Parasocial Interactions with Influencers in Spain. *Int. J. Adolesc. Youth* **2026**, *31*, 2670402. [CrossRef]

40. Salthouse, T.A. The Processing-Speed Theory of Adult Age Differences in Cognition. *Psychol. Rev.* **1996**, *103*, 403–428. [CrossRef] [PubMed]
41. Wang, L.; Yeap, J.A.L.; Liu, J.; Li, Z. From Avatars to Algorithms: Virtual Streamers and AI-Enabled Consumer Behavior in Live Streaming Commerce—A Systematic Review. *J. Theor. Appl. Electron. Commer. Res.* **2026**, *21*, 57. [CrossRef]
42. Mrad, M.; Ramadan, Z.; Nasr, L.I. Computer-Generated Influencers and Generational Responses to AI Disclosure. *Mark. Intell. Plan.* **2025**, *43*, 412–430.
43. Baek, T.H.; Kim, S.; Yoo, C.Y. Effect of Disclosing AI-Generated Content on Prosocial Advertising Evaluation. *Int. J. Advert.* **2026**, *45*, 288–310. [CrossRef]
44. Wojdynski, B.W.; Evans, N.J. Going Native: Effects of Disclosure Position and Language on the Recognition and Evaluation of Online Native Advertising. *J. Advert.* **2016**, *45*, 157–168. [CrossRef]
45. Boerman, S.C.; van Reijmersdal, E.A.; Neijens, P.C. Effects of Sponsorship Disclosure Timing on the Processing of Sponsored Content: A Study on the Effectiveness of European Disclosure Regulations. *Psychol. Mark.* **2017**, *31*, 214–224.
46. Arango Espinal, E.; Osorio Andrade, C.F.; Uribe Jiménez, M. Eficacia del Gobierno y su Relación con el Respaldo Ciudadano en Redes Sociales: El Caso del Paisaje Cultural Cafetero de Colombia. *Estud. Polít.* **2025**, *72*, 267–297. [CrossRef]
47. Coppock, A.; Guess, A.; Ternovski, J. When Treatments are Tweets: A Network Mobilization Experiment over Twitter. *Polit. Behav.* **2022**, *38*, 105–128.
48. Banker, S.; Park, J. Evaluating Prosocial COVID-19 Messaging Frames: Evidence from a Field Study on Facebook. *Judgm. Decis. Mak.* **2020**, *15*, 1037–1043. [CrossRef]
49. Gordon, B.R.; Zettelmeyer, F.; Bhargava, N.; Chapsky, D. A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook. *Mark. Sci.* **2023**, *38*, 193–225.
50. DANE. *Proyecciones de Poblacion Municipal 2018–2035*; Departamento Administrativo Nacional de Estadistica: Bogota, Colombia, 2024.
51. Rubin, D.B. Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment. *J. Am. Stat. Assoc.* **1980**, *75*, 591–593. [CrossRef]
52. Orazi, D.C.; Johnston, A.C. Running Field Experiments Using Facebook Split Test. *J. Bus. Res.* **2020**, *118*, 189–198. [CrossRef] [PubMed]
53. Kirk, C.P.; Peck, J.; Swain, S.D. The AI-Authorship Effect: Understanding Authenticity, Moral Disgust, and Consumer Responses to AI-Generated Marketing Communications. *J. Bus. Res.* **2025**, *178*, 114984. [CrossRef]
54. Banco de la República. Tasa de Cambio Representativa del Mercado (TRM). Available online: https://www.banrep.gov.co (accessed on 13 August 2026).
55. Rosenzweig, L.R.; Bergquist, P.; Hoffmann Pham, K.; Rampazzo, F.; Mildenberger, M. Survey Sampling in the Global South Using Facebook Advertisements. *Polit. Sci. Res. Methods* **2025**, *13*, 781–797. [CrossRef]
56. Ver Hoef, J.M.; Boveng, P.L. Quasi-Poisson vs. Negative Binomial Regression: How Should We Model Overdispersed Count Data? *Ecology* **2007**, *88*, 2766–2772. [CrossRef] [PubMed]
57. Becerra Rodríguez, F.; Vela Peón, F. Modelos de Conteo y sus Aplicaciones. In *Análisis de Datos con Stata*; Lara Rosano, F., Ed.; UNAM: Mexico City, Mexico, 2011; pp. 225–251.
58. Skrondal, A.; Rabe-Hesketh, S. *Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models*; Chapman & Hall/CRC: Boca Raton, FL, USA, 2004; Volume 4.
59. Chen, Y.; Zhang, Y.; Fan, H.; Guo, Y. Can AI Streamers Drive Purchases in Food Delivery E-Commerce Live Streaming? Evidence from an O2O Platform. *J. Retail. Consum. Serv.* **2026**, *90*, 104683. [CrossRef]
60. Hovland, C.I.; Janis, I.L.; Kelley, H.H. *Communication and Persuasion: Psychological Studies of Opinion Change*; Yale University Press: New Haven, CT, USA, 1953.
61. Breves, P.L.; Schramm, H. Bridging Psychological Distance: The Effect of Virtual Influencers on Consumers across Product Types. *Comput. Hum. Behav.* **2023**, *149*, 107927.
62. Forrai, M.; Balaban, D.C.; Schmuck, D. Disclosures and Literacy as Determinants of AI-Influencer Recognition and Well-Being. *Comput. Hum. Behav.* **2026**, *182*, 108978. [CrossRef]
63. Magni, M.; Maruping, L.M.; Hoegl, M.; Proserpio, L. Algorithm Aversion and Human Decision-Making: A Systematic Review. *Technol. Forecast. Soc. Chang.* **2023**, *187*, 122175.
64. Mahmud, H.; Islam, A.K.M.N.; Ahmed, S.I.; Smolander, K. What Influences Algorithmic Decision-Making? A Systematic Literature Review on Algorithm Aversion. *Technol. Forecast. Soc. Chang.* **2022**, *175*, 121390. [CrossRef]
65. Wu, L.; Dodoo, N.A.; Wen, T.J. Disclosing AI's Involvement in Advertising to Consumers: A Task-Dependent Perspective. *J. Advert.* **2025**, *54*, 20–38. [CrossRef]
66. Beckert, J.; Koch, T.; Viererbl, B.; Schulz-Knappe, C. The Disclosure Paradox: How Persuasion Knowledge Mediates Disclosure Effects in Sponsored Media Content. *Int. J. Advert.* **2021**, *40*, 1160–1186. [CrossRef]
67. Berlyne, D.E. *Conflict, Arousal, and Curiosity*; McGraw-Hill: New York, NY, USA, 1960.

68. Friestad, M.; Wright, P. The Persuasion Knowledge Model: How People Cope with Persuasion Attempts. *J. Consum. Res.* **1994**, *21*, 1–31. [CrossRef] [PubMed]
69. Liu, Q.; Ma, N.; Zhang, X. Can AI-Virtual Anchors Replace Human Internet Celebrities for Live Streaming Sales of Products? An Emotion Theory Perspective. *J. Retail. Consum. Serv.* **2025**, *82*, 104107. [CrossRef]

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