

From Novelty to Credibility? A Longitudinal Structural Topic Model of Audience Discourse towards Lil Miquela on YouTube

Nan Li

nan.li@sd.taylors.edu.my

Taylor's University

Luyao Yuan

UCSI University

Research Article

Keywords: virtual influencer, structural topic model, audience discourse, credibility, sponsorship, YouTube

Posted Date: September 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10691178/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Virtual influencers have moved from experimental digital characters to established marketing actors, but audience responses are usually examined at one point in time. This study investigates whether discourse around a mature virtual influencer shifted from ontological novelty towards credibility, authenticity, commercialisation or scepticism. It analyses public YouTube comments concerning Lil Miquela between January 2021 and July 2026. A census of 20 eligible official-channel videos and two verified brand uploads produced 3,353 unique comments and replies. The primary corpus comprised 1,778 English-language top-level comments after duplicate and vocabulary filtering. A 15-topic structural topic model estimated nonlinear temporal prevalence and sponsorship associations while controlling for comment likes and delay. Discourse did not follow a single novelty-to-credibility trajectory. Robot–human identity confusion declined by 6.2 percentage points, whereas music and creative-output discourse increased by 13.9 points. However, the largest increase concerned return and renewed attention (+ 23.3 points), showing that posting gaps and new releases strongly structured the longitudinal pattern. Sponsored-video comments contained more brand/product evaluation (+ 10.9 points), technological-evolution/advertising (+ 6.4 points) and credibility/corporate-control discourse (+ 2.4 points). The brand/product and credibility contrasts became inconclusive after uploader adjustment, while technological-evolution/advertising remained prominent. The findings reconceptualise virtual-influencer maturation as selective normalisation punctuated by content events and commercial context, rather than a uniform development of trust. They also demonstrate the value and limits of longitudinal STM for audience discourse when sponsorship support is sparse and platform activity is uneven.

1 Introduction

Virtual influencers (VIs) are digitally produced personae that publish content, interact with audiences and perform many of the communicative functions associated with human social-media influencers. Their appearance, biography and activity can be coordinated by creators or organisations, while the public-facing character is presented as a continuing social actor (Oliveira & Chimenti, 2021; Sands et al., 2022; Audrezet et al., 2025). This arrangement gives firms considerable control over the endorser, yet it also makes the conditions of authenticity unusually visible. A VI can appear emotionally expressive, socially available and aesthetically consistent while remaining a designed and managed commercial asset. Consequently, audience discourse around VIs is not simply a response to advertising. It is also a continuing negotiation of what the persona is, who or what acts through it, whether its expressions are credible, and whether a corporate construction can support a meaningful relationship.

Lil Miquela is an especially suitable case for examining this negotiation over time. Rather than representing a newly introduced character, she is a mature VI whose identity has been performed across music, fashion, personal storytelling and brand collaboration. Studies have used her to examine identity performance, authenticity, parasociality, sponsorship disclosure, sentiment and consumer responses (Lim & Lee, 2023; Koles et al., 2024; Looi et al., 2025; Elyamany et al., 2025). Nevertheless, much VI research still relies on experiments, surveys or cross-sectional samples. These designs identify

mechanisms such as human-likeness, credibility, social presence, empathy, product fit and disclosure, but they provide limited evidence about how naturally occurring audience discourse develops as the same virtual persona ages, changes content and re-enters commercial contexts (Belanche et al., 2024; Dabiran et al., 2024; Gutuleac et al., 2024; Sorosrungruang et al., 2024; Liu & Wang, 2025). Recent reviews accordingly identify longitudinal, platform-sensitive and context-sensitive research as a priority (Pushparaj et al., 2025; Khalfallah & Keller, 2025; Tian et al., 2026; Zhou et al., 2026; Hussain et al., 2026; Audrezet et al., 2025).

The longitudinal question matters because novelty is often invoked as both an advantage and a weakness of VIs. An unfamiliar synthetic persona can attract curiosity, yet attention based on ontological uncertainty may diminish as the format becomes recognised. This could allow audiences to evaluate creative output, credibility and brand fit rather than repeatedly asking whether the persona is human. It could equally expose commercial control, activate scepticism or make a VI less distinctive. Neither possibility implies a simple linear outcome. Credibility may develop in one part of the discourse while hostility, advertising awareness or identity confusion persists elsewhere. Posting gaps, releases and campaign events can also produce changes that resemble maturation but are actually responses to a different content mix.

Existing work already establishes that YouTube contains a substantial and analytically valuable record of audience responses to Lil Miquela. Feng et al. (2024) analysed 17,920 English comments and identified topics related to VI attributes, consumer trust and influencer–product fit. Choi et al. (2024) examined emotional reactions by content type and comment timing, finding different response profiles for music, narrative and external-channel videos. These studies demonstrate corpus viability, but their use of latent Dirichlet allocation or lexicon-based sentiment does not directly estimate how topic prevalence changes with continuous time and document metadata. In parallel, Looi et al. (2025) compared sponsored and non-sponsored Instagram posts and comments, revealing identity ambiguity and mixed emotions around sponsorship disclosure. That study concerns Instagram, a different audience environment, and does not model longitudinal topic prevalence. The present research therefore extends, rather than duplicates, this evidence.

The study determines how naturally occurring YouTube discourse around Lil Miquela changed from January 2021 to July 2026 and how topic prevalence was associated with sponsored versus non-sponsored videos. STM discovers latent themes while incorporating document-level covariates (Roberts et al., 2014; Roberts et al., 2019). Each comment is a document; time, sponsorship, likes and comment delay enter the primary specification, with uploader and corpus alternatives assessed in sensitivity models.

This research makes three contributions. First, it changes the empirical question from whether audiences generally accept VIs to how the composition of discourse around one established VI changes over time. This offers a more demanding test of the often-assumed novelty-to-credibility transition. Second, it separates sponsorship from uploader context. Because commercial videos may be posted by

the VI or by a brand, an apparent sponsorship difference may partly reflect the audience and communicative environment of the uploading channel. Third, it advances transparent text-as-data practice in VI research by reporting corpus attrition, topic-count diagnostics, uncertainty intervals, topic alignment and random-start sensitivity. The resulting account is deliberately associational: public comments reveal observable discourse within the sampled videos, not the stable attitudes of all viewers or a causal effect of sponsorship.

2 Conceptual background and research questions

2.1 Credibility, authenticity and the managed virtual source

Credibility provides one theoretical route through which a digital persona can become persuasive. Source-credibility research distinguishes trustworthiness, expertise and attractiveness as related but separable evaluations (Ohanian, 1990). In influencer marketing, credible communication and message value can strengthen trust in branded content, while influencer characteristics, follower characteristics, post design and platform context jointly shape effectiveness (Lou & Yuan, 2019; Belanche et al., 2021; Leung et al., 2022; Pan et al., 2025). For VIs, however, the source is simultaneously a social presentation and an engineered artefact. Expertise can be scripted, attractiveness can be rendered, and consistency can be centrally controlled. The traits that support credibility are therefore inseparable from knowledge that a human or corporate production system stands behind the persona.

Authenticity is similarly conditional. Human influencers manage a tension between commercial demands and a presentation of intrinsic motivation, personal continuity and audience intimacy (Audrezet et al., 2020; Zniva et al., 2023). VI research shows that authenticity can be manifested through coherent identity, transparent fiction, narrative consistency and meaningful interaction, rather than through literal human existence (Koles et al., 2024). Yet virtuality can also reduce perceived brand authenticity or make the endorser appear creepy, particularly when realism, social cues and audience expectations do not align (Gutuleac et al., 2024; Kim & Baek, 2024; Song et al., 2024). Comparative studies accordingly produce conditional rather than universal conclusions about whether human or virtual influencers are more effective (Kim et al., 2023; Belanche et al., 2024; Chu et al., 2026). The most recent syntheses locate credibility, authenticity, trust, interactivity and social relationality at the centre of VI outcomes but emphasise substantial moderation across contexts (Khalfallah & Keller, 2025; Zhou et al., 2026; Nguyen & Wu, 2026).

Credibility is also publicly negotiated when commenters debate whether a persona is believable, autonomous, commercially appropriate or suitable for an audience. A mature VI may become normalised without becoming trusted: commenters may cease asking what the character is while attending more to corporate programming, campaign fit or creative performance. Treating credibility as one endpoint would obscure these combinations.

2.2 Anthropomorphism, identity ambiguity and parasociality

Audience engagement with a VI depends partly on the application of human social knowledge to a nonhuman or ambiguously human entity. Anthropomorphism theory explains such attribution through elicited agent knowledge, the motivation to understand and predict an agent, and the desire for social connection (Epley et al., 2007). Human-likeness can support social presence and interaction, but near-human representation may also create discomfort when appearance or behaviour violates expectations (Mori et al., 2012). VI studies show that anthropomorphism, autonomy, emotional narrative and social cues affect credibility, parasocial response and willingness to engage (Lim & Lee, 2023; Dabiran et al., 2024; Gutuleac et al., 2024; Sorosrungruang et al., 2024; Liu & Wang, 2025). These effects are not merely visual. Commenters may discuss voice, animation, bodily possibility, consciousness, scripted interaction or the human operators behind the persona.

Parasocial interaction offers a relational interpretation of these responses. Originally used to describe an apparently intimate but structurally non-reciprocal relationship with a media performer (Horton & Wohl, 1956), the concept now informs research on influencers whose posts and comments create impressions of access and reciprocity. Audience comments can contribute to the perceived relationship among parasociality, credibility and endorsement response (Reinikainen et al., 2020). On YouTube, long-form address, recurring formats and creator vernacular can intensify the sense of relational continuity (Rohde & Mau, 2021; Munnukka et al., 2019). In a VI setting, however, parasociality can coexist with explicit awareness that interaction is staged or team-managed. Research comparing human and virtual sources links perceived similarity and human-likeness with parasocial interaction, while empathy can arise even when the agent is known to be virtual (Stein et al., 2024; Mirowska & Arsenyan, 2023).

Categorical confusion and social acceptance need not change together. Questions about whether Lil Miquela is a robot or human express ontological novelty, whereas talk about friendship, feelings or everyday routines attributes social and experiential capacity. Normalisation may therefore reduce identity confusion while creative or relational discourse increases, even if body and production questions persist.

2.3 Commercialisation, persuasion knowledge and sponsorship context

Commercial content makes the managed nature of influencer communication especially salient. The persuasion knowledge model proposes that people develop knowledge of persuasive agents and tactics and use it when interpreting persuasion attempts (Friestad & Wright, 1994). Sponsorship disclosures often increase advertising recognition and can activate persuasion knowledge, although their downstream effects depend on wording, context and existing attitudes (Boerman et al., 2017; Evans et al., 2017; Boerman, 2020; Karagür et al., 2022). Influencer credibility and disclosure can consequently have mixed relationships with purchase intention and engagement (Sesar et al., 2022). In VI marketing,

the issue is intensified because the influencer may itself be understood as a brand-owned or organisation-controlled communication vehicle.

VIs offer strategic advantages including consistency, design flexibility and reduced dependence on a human celebrity, but these advantages are inseparable from questions of falsity, accountability and control (Sands et al., 2022; Mouritzen et al., 2024). Product and audience congruence also matter: the effectiveness of an endorsement depends on how the influencer, product and consumer fit together (Belanche et al., 2021; De Veirman et al., 2017). An innovative technology campaign may make virtuality appear appropriate, whereas a low-fit collaboration may redirect discussion towards artificiality or commercial opportunism. The technology can therefore function as both the content of curiosity and a reason for perceived campaign fit.

The YouTube setting adds further context. Product promotion has become embedded in creator content, and visibility is shaped by platform conventions and audience expectations (Schwemmer & Ziewiecki, 2018; Bishop, 2019). A brand-channel upload is encountered within a different communicative environment from an upload on the VI's own channel. It may attract product-focused viewers, existing customers or technology enthusiasts who would not ordinarily comment on the persona. Therefore, sponsored versus non-sponsored and brand-uploaded versus VI-uploaded are not interchangeable classifications. Looi et al. (2025) found that Lil Miquela's sponsored and non-sponsored Instagram content produced different sentiment and engagement patterns, but also found persistent identity ambiguity. The present study tests whether sponsorship is associated with the prevalence of naturally discovered YouTube topics while explicitly assessing uploader confounding.

2.4 Longitudinal audience discourse

Most VI evidence estimates average effects, whereas a continuing persona operates through time. Recent work frames VIs as developing marketing actors whose authenticity, autonomy and cross-media activity require longitudinal and platform-specific research (Pushparaj et al., 2025; Tian et al., 2026; Hussain et al., 2026; Audrezet et al., 2025). Lil Miquela research has examined Instagram identity performance across several years and recommends analysing how audience trust and perceived reliability vary with post type and time (Elyamany et al., 2025). Theoretical reviews likewise argue that credibility, authenticity and engagement emerge through interacting design and contextual cues rather than a single static attribute (Khalfallah & Keller, 2025; Zhou et al., 2026; Nguyen & Wu, 2026).

A longitudinal corpus does not identify maturation automatically. Comments depend on posting dates, content, inactivity and whether a campaign reaches a different audience. Calendar time must be modelled flexibly and interpreted with video events because attention may rise around a release and reappear after absence. A sponsorship-time interaction also requires common support; later-only sponsored comments would confound sponsorship with period.

Based on these considerations, the study asks:

RQ1

What latent topics characterise English-language audience discourse towards Lil Miquela on YouTube?

RQ2

How did the prevalence of the discovered topics change between January 2021 and July 2026?

RQ3

How is topic prevalence associated with whether the relevant video is sponsored or non-sponsored?

RQ4

To what extent are sponsorship associations robust to uploader type and the prespecified engagement and timing controls?

RQ5

Within the 2023–2026 period of sponsorship overlap, did sponsorship-associated differences in topic prevalence vary over time?

3 Methodology

3.1 Research design and case boundary

The study uses a longitudinal, observational text-as-data design. Lil Miquela is treated as a theoretically informative single case rather than as a representative sample of all VIs. A single continuing persona holds the focal identity relatively stable while allowing discourse to be examined across calendar time, content events and commercial settings. This choice supports within-case temporal inference about topic prevalence, but it does not support statistical generalisation to other characters, platforms or cultural markets. The observation window was fixed as 1 January 2021 to 31 July 2026, providing complete calendar months and covering an active official-channel period, sponsored campaigns, an extended posting gap and renewed music-related activity.

The frame had two strata. A census of public official-channel videos yielded 20 eligible uploads. A second stratum retained the originating brand upload for each verified YouTube campaign featuring Lil Miquela. API searches and manual verification identified BMW (October 2023) and MSI (April 2024). Regional edits, dealer copies, portfolio reuploads, news, reactions, fan edits and incidental mentions were excluded. The final frame contained 22 videos; 21 comment sections were accessible and one official video had comments disabled.

This bounded census avoids selecting videos by popularity or comment volume. It also keeps the VI-channel and brand-channel environments observable rather than merging them. The frame cannot include deleted, private, regionally unavailable or undiscovered videos, and the brand-campaign census

is limited to campaigns with an eligible originating YouTube upload. These boundaries define the observable corpus and are not claims of total public discussion about Lil Miquela.

3.2 Comment retrieval and corpus construction

Public comment threads were collected through the official `commentThreads.list` endpoint of YouTube Data API v3, which supports video-specific pagination and returns published top-level threads (Google for Developers, 2026). Complete replies were collected through linked requests where available. Retrieval retained comment identifiers, video identifiers, publication times, like counts, parent relationships and plain text. Raw API responses were archived locally with extraction timestamps. The collection returned 3,353 unique comments and replies from the 21 accessible videos: 2,413 top-level comments (72.0%) and 940 replies (28.0%).

The primary unit of analysis was one top-level comment. Replies were excluded because they can respond primarily to another commenter rather than to the source video and are dependent within threads; they were retained for a sensitivity analysis. Comments had to be published within the fixed study window and classified as English by Lingua 2.1.1. The language policy yielded 2,014 English top-level comments. Their distribution was uneven: 1,302 were posted in 2021, 316 in 2022, 174 in 2023, 116 in 2024, 47 in 2025 and 59 in 2026. Of these comments, 1,844 were attached to non-sponsored videos and 170 to sponsored videos.

Exact comment identifiers were globally deduplicated during collection. A second content-level check normalised substantive text and retained the earliest observation when identical normalised content of at least 20 characters appeared more than once. This removed eight later copies and prevented repeated posting from receiving additional weight. User display names and author-channel identifiers were removed from processed analytical files. Because even de-identified public comments can be searchable, the study reports aggregate patterns and does not reproduce verbatim comments.

3.3 Video, sponsorship and document covariates

Sponsorship was coded at the video level from explicit commercial evidence. Evidence included originating brand-channel publication, partnership credits, campaign framing, product links or an official-channel brand collaboration. Videos were classified as sponsored only when the commercial relationship was documented; ordinary references to fashion, products or popular culture were insufficient. This produced three sponsored uploads in the final frame: the brand-uploaded BMW and MSI videos and Lil Miquela's official-channel MSI collaboration. Uploader type was coded separately as VI official channel or brand channel.

Calendar time was measured from each comment's publication date rather than from the video's publication date. `time_scaled` equalled days since 1 January 2021 divided by 365.25. Comment delay was the number of days between video publication and comment publication; its transformed value, `log1p_comment_delay`, controlled for the difference between immediate and later reactions.

log1p_comment_likes represented one plus the comment like count on a natural-log scale. Likes were included as an engagement-related control, not as an indicator of comment truth or audience representativeness. The primary model used sponsorship as the focal exposure and treated uploader as a sensitivity adjustment because uploader and sponsorship were strongly associated in the sparse campaign material.

3.4 Text preprocessing

Preprocessing was conducted in R as part of the STM workflow. Text was converted to lowercase, and punctuation, numbers and standard English stop words were removed using `stm::textProcessor`. Stemming was not applied because inflected and colloquial forms helped preserve meaning in short comments. Focal construct words such as *robot*, *human*, *real*, *fake*, *voice* and *feelings* were deliberately retained. Custom exclusions were limited to the focal name, basic YouTube/navigation tokens and web-address fragments, which could identify the sampled case without contributing thematic content.

Terms were required to occur in at least five documents. From 2,006 comments remaining after content deduplication, 228 became empty or contained only below-threshold terms and were removed. The final primary STM corpus contained 1,778 documents and 346 terms. It comprised 1,636 comments on non-sponsored videos and 142 on sponsored videos; 1,683 comments came from VI-channel uploads and 95 from brand-channel uploads. Figure 1 and Table 1 retain the complete attrition path so that topic prevalence is interpreted for the final analytical corpus rather than for all retrieved comments.

Table 1
Text preprocessing and corpus attrition

Corpus stage	Retained n	Excluded n	Terms	Rule or reason
All unique retrieved comments and replies in date window	3353	—	—	Official YouTube API collection from 21 accessible qualifying videos
Top-level comments only	2413	940	—	Replies excluded from the primary unit of analysis and retained for sensitivity analysis
English-language top-level comments	2014	399	—	Lingua 2.1.1 top detected language equals English
After later exact normalized-text copies removed	2006	8	—	Retain earliest occurrence when normalized substantive text has at least 20 characters
After STM text processing and minimum term-frequency rule	1778	228	346	Standard stop words, numbers and punctuation removed; no stemming; terms must occur in at least 5 documents; empty/rare-only documents dropped
Final STM vocabulary	—	—	346	Only terms surviving the document-frequency threshold
Final primary topic model	1778	0	346	K = 15 spectral STM with the frozen prevalence specification
<i>Note. The final primary model uses English top-level comments. Percentages and corpus characteristics are reported in the text and Fig. 1.</i>				

3.5 Structural topic model specification

STM models each document as a mixture of latent topics and each topic as a distribution over words, while permitting topic prevalence to vary with observed document metadata (Roberts et al., 2014). This makes it preferable to an LDA-only comparison when the research questions concern continuous time and sponsorship. All topic estimation, effect estimation, robustness models, correlation analysis and research figures were produced in R 4.6.1 using stm 1.3.8. Python supported API collection, language processing, privacy operations and non-graphical corpus preparation.

Candidate models were estimated for K = 8, 10, 12, 15, 18, 20 and 25. Selection considered held-out likelihood, residual dispersion, semantic coherence, exclusivity and representative-document interpretability rather than maximising a single statistic (Roberts et al., 2019; Grimmer & Stewart, 2013). Increasing K generally improved exclusivity and residual dispersion but reduced coherence and fragmented closely related discourse. K = 15 provided the strongest joint balance: exclusivity was 9.756,

held-out likelihood – 5.207 and residual dispersion 2.992, while its topics remained substantively distinguishable when high-probability comments were read. Figure 3 reports all diagnostics.

The primary $K = 15$ STM used spectral initialisation and the following prevalence formula:

$$\text{prevalence} = \sim s(\text{time_scaled}, \text{df} = 5) + \text{sponsored} + \log1p_comment_likes + \log1p_comment_delay.$$

A natural spline with five degrees of freedom allowed nonlinear temporal change without forcing a straight maturation trend. Topic-content covariates were not imposed, so sponsorship could change topic prevalence but not create group-specific word distributions. Effects were estimated with global uncertainty using 100 composition simulations in `estimateEffect`; plotted intervals used 1,000 coefficient simulations. Temporal predictions held other continuous covariates at their observed reference settings and reported expected proportions over the supported comment-date range. Sponsorship contrasts were adjusted sponsored-minus-non-sponsored differences in expected topic proportion. Differences are expressed in percentage points with 95% simulation intervals. Intervals spanning zero are described as inconclusive rather than as evidence of no association.

3.6 Topic interpretation and validation

Topic labels are interpretive summaries, not estimated psychological constructs. Each topic was examined using highest-probability words, FREX words, lift and score metrics, followed by close reading of ten high-probability comments. The four lexical views help distinguish frequent words from words that are unusually exclusive to a topic. Labels were kept data-near and combined multiple elements where representative comments contained a coherent mixture, for example, “credibility, corporate control and audience fit”. Theory mapping occurred after this lexical and documentary review, reducing the risk of forcing the output into preselected constructs (Grimmer & Stewart, 2013; Hannigan et al., 2019).

Validation followed a documented audit trail consistent with recommendations for text-as-data research (Birkenmaier et al., 2024). However, no independent second coder completed the blinded topic-validation workbook before manuscript drafting. The study therefore does not report inter-coder agreement or call the labels independently validated. Labelling was supported by multi-metric lexical evidence, representative-document review and comparison with the VI literature, with the absence of independent coding retained as a limitation. This disclosure is preferable to treating apparently intuitive labels as model facts.

3.7 Supporting and robustness analyses

Topic correlations were estimated with `stm::topicCorr`. The resulting network describes topics that tend to co-occur within documents after model estimation; it is neither a causal pathway nor a significance graph. A threshold of 0.05 was used for plotted positive edges, and Louvain communities assisted descriptive interpretation.

Four main sensitivity specifications assessed RQ4. First, an uploader-adjusted prevalence model added brand versus VI uploader. Second, a high-confidence-English model retained comments with language

confidence of at least .50. Third, an official-channel-only model removed the brand uploads. Fourth, a reply-inclusive model added replies and controlled for comment role; it was treated as a corpus-inclusion check because replies remain thread-dependent. Independently fitted topics were aligned to the primary solution by maximum one-to-one cosine similarity over common topic-word vocabularies. Weak alignment limits substantive comparison even when a coefficient appears similar.

Random-start stability was evaluated by aligning three additional $K = 15$ solutions to the spectral model. Stability was summarised through topic-word L1 distance on terms covering 90% of reference probability mass, top-ten word overlap and top-document overlap (Marani & Baumer, 2024). Finally, RQ5 was investigated only in the 2023–2026 overlap period, using 319 documents and a linear $\text{time_scaled} \times \text{sponsored}$ specification with the same controls. This model was retained in the appendix because its mean topic alignment with the primary model was weak. Estimating that interaction over 2021–2026 would be unsupported: there were no sponsored comments in 2021–2022.

3.8 Ethical and data-governance approach

The study used public, non-interactive YouTube data and did not contact commenters or access login-only information. Nevertheless, public availability was not treated as sufficient ethical justification. Context-sensitive internet-research guidance emphasises participant expectations, traceability and the lifecycle of collected data (Kinder-Kurlanda & Weller, 2020; Association of Internet Researchers, 2020). Accordingly, processed files exclude usernames and author identifiers, representative comments remain in restricted local storage, and no searchable quotation appears in the article. Raw text is not proposed for public deposit. Reproducibility is instead supported through collection and analysis code, video-level eligibility records, model specifications, aggregate tables and R figure scripts. The corresponding author must obtain or document the applicable institutional ethics determination before submission.

4 Results

4.1 Topic structure of audience discourse

RQ1 asked what topics characterised the corpus. The $K = 15$ solution revealed a discourse organised around four broad but overlapping areas: commercial evaluation; fandom and creative output; embodiment, agency and identity; and relationship, resistance and temporal re-engagement. Table 2 reports every topic's mean prevalence, defining words and interpretive scope.

The most prevalent topic was robot–human identity confusion (Topic 11, 10.18%). Its words included *robot*, *human*, *know*, *confused* and *program*, and representative documents directly questioned the persona's ontological status. Interaction performance and agency (Topic 10, 8.59%) concerned whether comments, social encounters and apparently spontaneous behaviour were scripted or enacted by others. Fandom and pop-culture affiliation (Topic 2, 7.99%) captured affectionate address, fan identity and references to fashion and popular performers. Production mechanism and voice (Topic 4, 7.55%)

and face, body and physical embodiment (Topic 9, 7.49%) addressed CGI, voice acting, face editing, physical presence and the human production system.

Other topics moved beyond ontological questions. Music and creative output (Topic 3, 7.20%) evaluated songs, releases and artistic work. Parasocial friendship and embodied routines (Topic 5, 6.13%) combined friendship language with questions about food, everyday activity and collaboration. Feelings, consciousness and empathy (Topic 8, 4.83%) debated whether the persona could care, feel or possess a mind. Robot sociality and conversational curiosity (Topic 15, 4.19%) concerned talking with robots and the possibility of relational acceptance. Together, these topics show that humanisation was negotiated through agency, embodiment, experience and interaction rather than through visual realism alone.

Commercial discourse was also differentiated. Brand and product evaluation (Topic 1, 5.66%) focused on BMW, MSI, product qualities and campaign execution. Technological evolution and advertising (Topic 6, 6.02%) linked improved realism and the future of virtual production with advertising use. Credibility, corporate control and audience fit (Topic 12, 6.16%) questioned seriousness, programming, corporate motives and what an intended audience might accept. Scepticism, cringe and social legitimacy (Topic 14, 5.66%) more directly rejected or defended the legitimacy of supporting a VI. Hostility, absence and AI threat (Topic 7, 5.77%) combined hostile or fearful reactions with demands that the inactive influencer return, indicating that resistance and audience investment could occur in the same topic. Return and renewed attention (Topic 13, 6.59%) captured welcome-back reactions and renewed release interest after inactivity.

Table 2
Topic summary and labelling evidence

Topic	Data-near label	Mean %	Highest-probability words	FREX words	Interpretive scope
1	Brand and product evaluation	5.66	one, look, hey, bmw, first, question, guys, car, nice, bad, wrong, tho	hey, bmw, question, bad, look, guys, one, last, trust, wrong, anyway, car	Discussion of BMW/MSI products campaign execution and product quality
2	Fandom and pop-culture affiliation	7.99	love, omg, cute, blackpink, fan, hair, miq, queen, lisa, pls, doesnt, lmao	love, omg, blackpink, fan, miq, queen, lisa, cute, lmao, fans, god, pls	Affection fan identity fashion and BLACKPINK/Lisa references
3	Music and creative output	7.20	much, music, good, make, new, need, song, better, thats, songs, great, got	music, new, got, song, make, good, thats, songs, whats, work, much, robotic	Requests and evaluations concerning songs releases and creative work
4	Production mechanism and voice	7.55	just, voice, wait, also, actually, amazing, animation, computer, watching, stop, generated, wanna	voice, actually, watching, animation, wait, actor, started, model, bro, computer, just, wonder	Questions or claims about animation CGI voice acting and how the persona is produced
5	Parasocial friendship and embodied routines	6.13	best, please, cool, friend, pretty, never, way, early, friends, tell, stupid, eat	best, cool, friend, pretty, early, friends, please, tell, eat, food, okay, way	Friendship collaboration and questions about eating or everyday embodied activity
6	Technological evolution and advertising	6.02	looks, now, wow, year, getting, thought, sounds, creepy, many, guess, ago, commercial	looks, getting, thought, creepy, now, guess, ago, commercial, future, using, didnt, wow	Comments on improving realism commercial use and future development
7	Hostility absence and AI threat	5.77	girl, get, life, will, thing, going, come, hate, little, scary, miss, whole	hate, little, will, get, thing, whole, fucking, life, scary, come, girl, reply	Hostile or fearful reactions plus demands that the absent influencer return

Note. Labels were supported by four word metrics and close reading of ten representative comments per topic. No independent-coder agreement is claimed.

Topic	Data-near label	Mean %	Highest-probability words	FREX words	Interpretive scope
8	Feelings consciousness and empathy	4.83	really, cant, believe, ive, channel, talking, kind, feelings, happy, remember, yall, mind	really, cant, believe, talking, feelings, anyone, seen, ive, happy, care, else, yall	Debate about whether the persona can feel care or possess consciousness
9	Face body and physical embodiment	7.49	real, people, world, face, person, body, take, everything, character, saw, mean, cuz	real, face, person, body, world, people, saw, head, screen, eyes, everything, mean	Explanations of the real body edited face and physical-world embodiment
10	Interaction performance and agency	8.59	like, can, actual, humans, say, comment, literally, feels, game, every, act, show	comment, actual, feels, can, humans, like, act, game, say, damn, literally, editing	Questions about real-world interaction comment engagement scripts and staged relationships
11	Robot-human identity confusion	10.18	robot, know, human, even, still, confused, someone, yes, crazy, water, happens, program	know, robot, confused, someone, water, happens, human, even, crazy, yes, still, program	Direct uncertainty about whether Miquela is a robot human or virtual entity
12	Credibility corporate control and audience fit	6.16	see, want, made, feel, don't, right, kinda, sorry, things, perfect, gonna, isnt	right, want, kinda, feel, see, isnt, maybe, don't, things, doesn't, irl, made	Scepticism about seriousness corporate motives target audience programming and what sells
13	Return and renewed attention	6.59	youre, beautiful, wtf, animated, hope, 're, back, keep, something, shit, years, missed	youre, beautiful, wtf, hope, 're, back, shit, missed, old, animated, part, keep	Welcome-back reactions after periods of absence and renewed releases
14	Scepticism cringe and social legitimacy	5.66	dont, think, shes, time, lol, weird, meet, give, exist, support, cringe, everyone	shes, think, weird, dont, support, cringe, everyone, give, time, click, lol, probably	Rejection warnings disbelief and arguments over whether virtual influencers deserve support

Note. Labels were supported by four word metrics and close reading of ten representative comments per topic. No independent-coder agreement is claimed.

Topic	Data-near label	Mean %	Highest-probability words	FREX words	Interpretive scope
15	Robot sociality and conversational curiosity	4.19	robots, ever, well, talk, said, wish, always, name, follow, speak, yeah, day	robots, talk, name, said, wish, speak, well, follow, always, questions, ever, yeah	Desire to converse with robots questions about their experience and relational acceptance
<i>Note. Labels were supported by four word metrics and close reading of ten representative comments per topic. No independent-coder agreement is claimed.</i>					

The correlation network in Fig. 4 reinforces these distinctions while showing their connections. The strongest relationship joined interaction performance/agency and robot–human identity confusion ($r = .617$). Identity confusion was also related to scepticism/social legitimacy ($r = .458$) and face/body embodiment ($r = .422$). Brand/product evaluation correlated with technological evolution/advertising ($r = .396$), while fandom correlated with parasocial friendship ($r = .374$). Music/creative output correlated with return/renewed attention ($r = .340$). These patterns suggest that the corpus contains an ontology-and-agency cluster, an affiliation cluster and a commercial-technology cluster. Correlations indicate co-occurrence, not an ordered process.

4.2 Longitudinal changes in topic prevalence

RQ2 concerned change over the supported prediction range, from 23 February 2021 to 29 July 2026. The results do not support a general movement from novelty to credibility. Instead, several ontology- and interaction-oriented topics declined while creative work and event-specific re-engagement became more prominent. Robot–human identity confusion decreased from an estimated 11.06% to 4.89%, a change of – 6.17 percentage points. Interaction performance and agency declined by 5.09 points, parasocial friendship and embodied routines by 4.73 points, feelings/consciousness/empathy by 3.76 points, robot sociality/conversational curiosity by 3.39 points and production mechanism/voice by 3.09 points. Face/body embodiment declined more modestly by 2.11 points.

Two increases were much larger. Music and creative output rose from 4.84% to 18.77% (+ 13.93 points). Return and renewed attention rose from 0.41% to 23.67% (+ 23.26 points). Figure 5 shows that these changes were not smooth substitutions. The return topic was concentrated around the reappearance of official uploads and a 2026 music release after a long gap. The creative-output increase similarly corresponded with new song-related material. Fandom and pop-culture affiliation fell by 8.20 points, which further indicates that later discourse did not simply become more favourable. Rather, its referent changed from general fandom and identity questions towards the new creative event.

The focal commercial and evaluative topics showed smaller end-to-start changes. Technological evolution and advertising increased by 1.54 points and brand/product evaluation by 0.98 points. Hostility, absence and AI threat increased by 1.57 points. By contrast, credibility, corporate control and

audience fit declined by 2.98 points, and scepticism, cringe and social legitimacy declined by 1.73 points. Thus, explicit credibility and scepticism did not rise as dominant endpoints. Their presence remained substantively meaningful, but the largest longitudinal changes concerned reduced identity ambiguity and renewed attention to creative output.

Figure 6 places these curves beside the video schedule. Most qualifying official uploads occurred in 2021, sponsored campaigns appeared in 2023–2024, and official uploads resumed in 2026. The sparse later-year counts and content discontinuities mean that predicted curves summarise the observed comment stream, not a balanced panel of identical content. The apparent maturation process is therefore better described as selective normalisation punctuated by upload events.

4.3 Sponsorship associations

RQ3 asked whether topics differed between comments on sponsored and non-sponsored videos. Seven adjusted contrasts had 95% intervals that did not span zero. Brand and product evaluation was 10.90 percentage points more prevalent for sponsored videos (95% interval [8.67, 13.46]). Technological evolution and advertising was 6.38 points higher [3.69, 9.22], and credibility, corporate control and audience fit was 2.40 points higher [0.70, 4.29]. These results indicate that sponsored contexts redirected discourse towards the product, the commercial use of virtual technology and questions about the controlled source.

Four topics were less prevalent in sponsored-video comments. Return and renewed attention differed by – 5.48 points [– 8.80, – 2.09], fandom and pop-culture affiliation by – 4.80 [– 7.94, – 1.67], interaction performance and agency by – 2.11 [– 3.88, – 0.39], and scepticism/cringe/social legitimacy by – 1.86 [– 3.57, – 0.14]. The reduced scepticism topic should not be read as greater trust: sponsored discourse simultaneously contained more corporate-control and technology-as-advertising discussion. The pattern is one of issue redirection rather than a unidimensional positive or negative response.

Intervals for the other eight topics spanned zero. In particular, the – 1.58-point identity-confusion difference [– 3.26, 0.11] was inconclusive, consistent with evidence that ontological ambiguity can remain salient across sponsored and non-sponsored material (Looi et al., 2025). The full set of contrasts appears in Table 3 and Fig. 7.

Table 3
Primary longitudinal and sponsorship estimates

Topic	Label	Start %	End %	Change (pp)	Sponsored – non-sponsored pp [95% interval]
1	Brand and product evaluation	2.07	3.05	0.98	10.90 [8.67, 13.46]
2	Fandom and pop-culture affiliation	15.55	7.34	-8.20	-4.80 [-7.94, -1.67]
3	Music and creative output	4.84	18.77	13.93	1.23 [-2.37, 4.81]
4	Production mechanism and voice	6.66	3.57	-3.09	-0.71 [-3.23, 1.80]
5	Parasocial friendship and embodied routines	6.54	1.81	-4.73	-1.30 [-3.67, 0.95]
6	Technological evolution and advertising	2.74	4.28	1.54	6.38 [3.69, 9.22]
7	Hostility absence and AI threat	5.02	6.58	1.57	1.80 [-0.36, 4.03]
8	Feelings consciousness and empathy	7.32	3.55	-3.76	-1.58 [-3.26, 0.14]
9	Face body and physical embodiment	8.40	6.29	-2.11	-1.54 [-4.28, 1.01]
10	Interaction performance and agency	9.33	4.24	-5.09	-2.11 [-3.88, -0.39]
11	Robot-human identity confusion	11.06	4.89	-6.17	-1.58 [-3.26, 0.11]
12	Credibility corporate control and audience fit	7.01	4.03	-2.98	2.40 [0.70, 4.29]
13	Return and renewed attention	0.41	23.67	23.26	-5.48 [-8.80, -2.09]
14	Scepticism cringe and social legitimacy	6.23	4.50	-1.73	-1.86 [-3.57, -0.14]
15	Robot sociality and conversational curiosity	6.82	3.43	-3.39	-1.73 [-3.70, 0.27]
<i>Note. Start and end are 23 February 2021 and 29 July 2026, the supported prediction range. All differences are percentage points; sponsorship estimates are adjusted associations.</i>					

4.4 Robustness, uploader context and exploratory interaction

RQ4 required the sponsorship interpretation to survive alternative specifications. Adding uploader type materially changed the results despite near-perfect topic alignment with the primary model (mean cosine similarity = .992). Brand/product evaluation changed from + 10.90 to - 1.29 points [- 3.71, 1.20], and credibility/corporate control changed from + 2.40 to + 0.32 [- 1.82, 2.45]. Return/renewed attention and scepticism also became inconclusive. These changes show that their primary sponsorship differences were entangled with brand-channel context. A product-focused audience on a BMW or MSI channel is not equivalent to the audience of the VI's own uploads.

Technological evolution and advertising remained higher in the uploader-adjusted model and increased to + 11.99 points [7.26, 16.28]. Hostility, absence and AI threat also became higher (+ 7.37 [3.29, 11.45]), while fandom remained lower (- 4.03 [- 8.01, - 0.03]). The persistence of the technology topic is plausible because the campaigns themselves presented a synthetic spokesperson alongside technologically innovative products. It remains an adjusted association, not evidence that sponsorship caused technological optimism.

Other corpus sensitivities were less directly comparable. The high-confidence-English model contained 1,407 documents, the official-channel-only model 1,683, and the reply-inclusive model 2,420. Their mean topic alignments were .480, .456 and .436 respectively. Figure 8 therefore treats them as structural checks and visually marks alignment rather than implying exact replication. The official-only comparison also contained only 47 sponsored documents from a single campaign, which substantially limits generality. Replies added conversational material but remain nested and cannot be treated as independent confirmations.

Random initialisations further showed that the boundaries among short-comment topics were solution-dependent. Across three aligned random-start solutions, mean L1 distance was 0.684, mean overlap among the top ten words was 1.71 words, and mean top-document overlap was 0.578. The spectral solution was reproducible as the declared primary model, but alternative initialisations partitioned some vocabulary differently. Accordingly, confidence should be placed in broad, repeatedly observable discourse families and in estimates conditional on the frozen solution, not in the inevitability of every fine-grained label.

Table 4
Robustness and stability summary

Specification	Documents	Sponsored	Mean alignment	Result or limit	Interpretation
Primary	1778	142	1.000	7 topic contrasts had 95% intervals excluding zero	Primary adjusted association
Uploader adjusted	1778	142	0.992	3 topic contrasts had 95% intervals excluding zero; uploader materially changes focal estimates	Uploader confounding prevents causal sponsorship interpretation
English confidence $\geq .50$	1407	119	0.480	6 topic contrasts had 95% intervals excluding zero	Language classification uncertainty sensitivity
Official channel only	1683	47	0.456	3 topic contrasts had 95% intervals excluding zero; only 47 sponsored documents	Sparse single-campaign sponsored comparison
Replies included	2420	194	0.436	6 topic contrasts had 95% intervals excluding zero; replies are thread-dependent	Corpus-inclusion sensitivity, not an independent-observation analysis
2023–2026 interaction	319	139	0.352	Appendix-only because alignment is weak and full-window sponsorship support is absent	Exploratory structural sensitivity only
Random starts	1778	142	—	Mean aligned L1 distance 0.6843; mean top-10 word overlap 1.7111	Topic boundaries are solution-dependent
<i>Note. Alignment is cosine similarity after maximum one-to-one matching with the primary topic-word solution. Random-start entries report topic-boundary stability rather than sponsorship estimates.</i>					

RQ5 was examined in the 319-document 2023–2026 overlap sample. Mean alignment with the primary model was only .352. The time-by-sponsorship coefficients are therefore retained as an exploratory appendix figure and do not support topic-specific longitudinal claims. This result is substantively important: the available data cannot identify whether sponsorship differences themselves evolved across the full period. The appropriate conclusion is limited to average adjusted associations within the observed content mix.

5 Discussion

5.1 Selective normalisation, not a linear credibility transition

The study began with a deliberately open question: has mature VI discourse moved from novelty towards credibility? The answer is qualified. The clearest novelty marker, robot–human identity confusion, became less prevalent, as did discourse about production, embodiment, feelings and staged interaction. Audiences were less occupied at the end of the observed period with resolving what Lil Miquela was. This is consistent with normalisation of the category and with contemporary definitions that place VIs within an established influencer system rather than treating them solely as technological curiosities (Tian et al., 2026; Audrezet et al., 2025).

However, declining ontological uncertainty did not produce a parallel rise in credibility discourse. Credibility/corporate control and social-legitimacy scepticism both declined slightly over the full range, while the dominant later patterns concerned a return and a song release. The results therefore separate familiarity from trust. Audiences can become less confused about the category without regarding the persona as more credible, authentic or socially legitimate. This distinction refines work that treats credibility as a relatively stable response to human-likeness, disclosure or source attributes (Lim & Lee, 2023; Dabiran et al., 2024; Chu et al., 2026). In an ongoing audience environment, the topic occupying attention may depend more immediately on what the persona has released and how long it has been absent.

The finding also qualifies the novelty argument in VI marketing. Novelty did not simply disappear; it changed form. Direct identity confusion declined, but technological-evolution and advertising discourse remained, particularly in sponsored contexts. Virtuality became less a categorical puzzle and more an applied commercial technology. This supports conceptualisations of VIs as both social sources and configurable marketing artefacts (Sands et al., 2022; Mouritzen et al., 2024). The normalisation observed here is therefore selective: ontology receded, while technological and corporate mediation remained available for evaluation.

5.2 Event time and the meaning of longitudinal change

The largest result was the growth of return/renewed attention, followed by music/creative output. Without the event timeline, these increases could be misinterpreted as a gradual movement towards creative legitimacy. Instead, the upload record shows a long inactive period followed by 2026 music-related posts. The longitudinal model correctly identifies a change in discourse prevalence, but substantive interpretation requires the content history. Calendar time in platform research is partly event time.

This point contributes to calls for longitudinal VI research (Pushparaj et al., 2025; Zhou et al., 2026; Elyamany et al., 2025). Following one persona over years does not by itself distinguish ageing,

technological development, audience learning and changes in the content supplied. Researchers should therefore combine flexible time effects with an explicit event inventory and common-support assessment. A balanced repeated-content design would answer a different question, but it is rarely available for naturally occurring influencer activity. In such settings, the defensible estimand is how discourse in the observed stream changed, not how the same audience would react to an identical stimulus after five years.

The event-centred interpretation is also consistent with YouTube's creator environment. Visibility and participation are organised through upload rhythms, platform imaginaries and creator–audience conventions (Schwemmer & Ziewiecki, 2018; Bishop, 2019). When a channel returns, comments can express memory, absence and re-entry before they evaluate the persona's stable qualities. “Maturity” should thus be conceptualised as a layered process that includes category familiarity, relationship continuity and episodic salience. This layered view also prevents a burst of supportive attention after absence from being treated as evidence that sceptical audiences have converted into trusting followers.

5.3 Sponsorship changes what audiences discuss

Sponsored videos were associated with more product evaluation, technological-advertising discourse and corporate-control/fit discourse. This pattern extends sponsorship-disclosure research by showing topic displacement rather than only advertising recognition, sentiment or purchase response (Boerman et al., 2017; Evans et al., 2017; Karagür et al., 2022; Looi et al., 2025). Commercial context made the object of persuasion and the utility of the virtual source more explicit. It did not produce a uniformly sceptical corpus: the direct scepticism topic was slightly lower in the primary model, while corporate-control questions were higher.

Persuasion knowledge helps explain this coexistence. Recognition of commercial intent can direct attention towards who controls the message, whether the persona fits the product and how the campaign is executed (Friestad & Wright, 1994). Yet disclosure and sponsorship do not determine one evaluative direction. Transparent or well-fitted commercial content can be accepted, while a synthetic endorser can still trigger questions about control (Boerman, 2020; Sesar et al., 2022). This is especially relevant for technology products, where the VI may signal innovation and campaign congruence (Belanche et al., 2021; Kim et al., 2023).

Uploader adjustment substantially weakened the product and credibility contrasts. This is not merely a statistical inconvenience; it identifies a managerial boundary. A brand upload enters a channel with its own audience, norms and product interests. Brands cannot assume that comment discourse observed under a campaign reflects a changed attitude among the VI's usual followers. The robust technology/advertising topic suggests that the synthetic format can focus attention on innovation, but the effect is inseparable from campaign category and placement. Future studies need multiple campaigns within each uploader type or matched cross-posts before making stronger claims about sponsorship.

5.4 Theoretical and managerial implications

Theoretically, the findings support a compositional view of VI credibility. Trustworthiness, expertise and attractiveness remain useful dimensions (Ohanian, 1990), but naturally occurring discourse distributes credibility-related judgement across identity, agency, production, commercial control, product fit and creative performance. Authenticity is similarly plural: a VI may be accepted as a coherent fictional character while its organisational motives remain visible (Audrezet et al., 2020; Zniva et al., 2023; Koles et al., 2024). Topic modelling makes this plurality observable without requiring every comment to be classified as authentic or inauthentic.

The ontology-and-agency correlation cluster also links anthropomorphism with accountability. Strong co-occurrence among identity confusion, embodiment, interaction performance and scepticism indicates that human-like appearance is only one part of the audience problem. Commenters also seek an account of who speaks, who acts and how social behaviour is produced. This complements evidence that anthropomorphism and social cues can increase connection or creepiness depending on design and disposition (Epley et al., 2007; Gutuleac et al., 2024; Kim & Baek, 2024; Liu & Wang, 2025). For a managed digital persona, perceived agency may be as consequential as perceived humanness.

Managerially, campaign evaluation should distinguish favourable attention from topic relevance. More discussion of technology or products can indicate campaign salience, but it does not establish trust or purchase intention. Brands should compare audience discourse across both the VI's channel and the brand channel, monitor corporate-control and fit themes, and avoid interpreting a low frequency of explicit scepticism as absence of persuasion knowledge. Creative continuity also matters. The strong return topic suggests that inactivity changes the meaning of renewed engagement; a relaunch may need to rebuild narrative continuity before commercial metrics can be compared with earlier periods.

5.5 Limitations and future research

The first limitation is case and platform specificity. Lil Miquela is a highly established, human-like VI, and YouTube supports different content forms and comment norms from Instagram or TikTok. Cross-case research should compare human-like, anime-like, brand-owned and more autonomous VIs across platforms. Second, the study observes comments, not all viewers. Commenters self-select, and public comments cannot establish demographic composition, private attitudes or purchasing behaviour. Like counts were used only as a control.

Third, the temporal corpus is unbalanced. Most comments appeared in 2021, sponsored material began in 2023, and official activity resumed in 2026. The spline describes this observed support but cannot isolate maturation from every content change. Future research could use matched content categories, interrupted time-series designs around identified events or recurring collection waves. Fourth, only three sponsored videos and two external brand campaigns qualified. Uploader was therefore strongly confounded with sponsorship. The uploader-adjusted model exposes this limitation but cannot create missing comparison cells.

Fifth, topic models are conditional representations. Short comments, a 346-term vocabulary and alternative initialisations produced solution dependence. The paper reports this rather than presenting $K = 15$ as uniquely true. Topic alignment and random-start diagnostics should become routine in longitudinal marketing text research (Birkenmaier et al., 2024; Marani & Baumer, 2024). Sixth, topic labels were not independently coded. Multi-metric words and representative documents provided an auditable basis, but future revision should include blinded human validation and adjudication. Finally, the analysis excludes unavailable, deleted and non-English discourse and protects raw comment text from public release. Cross-lingual work and ethically governed rehydration designs would broaden coverage while respecting contextual privacy.

6 Conclusion

This study used a longitudinal STM of 1,778 English-language top-level YouTube comments to examine how discourse around Lil Miquela changed from 2021 to 2026. The evidence does not support a simple transition from novelty to credibility. Identity confusion and several embodiment or agency topics declined, which suggests growing category familiarity. Yet credibility did not become the dominant destination. Later discourse was structured most strongly by renewed activity and creative output, while virtuality remained commercially salient as technology and advertising.

Sponsored contexts redirected comments towards products, technological development and corporate control, but uploader adjustment removed the apparent brand/product and credibility differences. The durable result was not a causal sponsorship effect; it was the importance of distinguishing commercial status from channel context. Overall, mature VI discourse is best understood as selective normalisation punctuated by content events and mediated by where commercial communication appears. For research, this requires models that join latent content with time, metadata, event histories and stability checks. For practice, it means that familiarity with a virtual persona should not be mistaken for credibility, and campaign attention should not be mistaken for audience trust.

Declarations

Data availability

Collection and analysis code, the video log, aggregate tables, model specifications and R figure scripts are available in the project package. Raw comment text is not publicly deposited because it is searchable user-generated content that may be edited or deleted. Any restricted access must follow the applicable institutional process and YouTube terms.

Ethical considerations.

The study used public, non-interactive YouTube comments, collected no private account data, removed user identifiers from processed files and reports no verbatim comments. Ethical approval has been granted by the ethical committee of the university.

Funding, competing interests and author contributions.

The authors did not receive any fundings from any institutions, organizations, or not-for-profit organizations. The authors declare no competing interests.

Use of generative AI.

Generative artificial intelligence tools were not utilised for the conception, data analysis, drafting or development of this research manuscript. Generative AI was only employed to identify grammatical errors and refine linguistic accuracy in the text, with all substantive content, arguments and interpretations remaining the original work of the authors.

Author Contribution

All authors made substantial contributions to the conceptualization of the work. Li Nan: Data Collection; analysis; methodology; writing original draft. Yuan Luyao: Data collection; data analysis; review and editing.

References

1. Association of Internet Researchers (2020). Internet research: Ethical guidelines 3.0. <https://aoir.org/reports/ethics3.pdf>
2. Audrezet, A., Kerviler, G. d., & Moulard, J. G. (2020). Authenticity under threat: When social media influencers need to go beyond self-presentation. *Journal of Business Research*, 117, 557–569. <https://doi.org/10.1016/j.jbusres.2018.07.008>
3. Audrezet, A., Koles, B., Moulard, J. G., Ameen, N., & McKenna, B. (2025). Virtual influencers: Definition and future research directions. *Journal of Business Research*, 200, 115647. <https://doi.org/10.1016/j.jbusres.2025.115647>
4. Belanche, D., Casaló, L. V., Flavián, M., & Ibáñez-Sánchez, S. (2021). Building influencers' credibility on Instagram: Effects on followers' attitudes and behavioral responses toward the influencer. *Journal of Retailing and Consumer Services*, 61, 102585. <https://doi.org/10.1016/j.jretconser.2021.102585>
5. Belanche, D., Casaló, L. V., Flavián, M., & Ibáñez-Sánchez, S. (2021). Understanding influencer marketing: The role of congruence between influencers, products and consumers. *Journal of Business Research*, 132, 186–195. <https://doi.org/10.1016/j.jbusres.2021.03.067>
6. Belanche, D., Casaló, L. V., & Flavián, M. (2024). Human versus virtual influences, a comparative study. *Journal of Business Research*, 173, 114493. <https://doi.org/10.1016/j.jbusres.2023.114493>
7. Birkenmaier, L., Lechner, C. M., & Wagner, C. (2024). The search for solid ground in text as data: A systematic review of validation practices and practical recommendations for validation.

- Communication Methods and Measures*, 18(3), 249–277.
<https://doi.org/10.1080/19312458.2023.2285765>
8. Bishop, S. (2019). Managing visibility on YouTube through algorithmic gossip. *New Media & Society*, 21(11–12), 2589–2606. <https://doi.org/10.1177/1461444819854731>
 9. Boerman, S. C. (2020). The effects of the standardized Instagram disclosure for micro- and meso-influencers. *Computers in Human Behavior*, 103, 199–207.
<https://doi.org/10.1016/j.chb.2019.09.015>
 10. Boerman, S. C., Willemsen, L. M., & Van Der Aa, E. P. (2017). ‘This post is sponsored’: Effects of sponsorship disclosure on persuasion knowledge and electronic word of mouth in the context of Facebook. *Journal of Interactive Marketing*, 38(1), 82–92.
<https://doi.org/10.1016/j.intmar.2016.12.002>
 11. Choi, M., Oh, Y. K., Cheong, H. R., & Chung, D. (2024). Emotional response to virtual influencer content on YouTube. *Journal of Digital Contents Society*, 25(9), 2563–2574.
<https://doi.org/10.9728/dcs.2024.25.9.2563>
 12. Chu, T. H., Looi, J., & Liu, D. (2026). A meta-analytic comparison of human influencers and virtual influencers across different levels of human-likeness: Is humanness more persuasive?. *Journal of Advertising Research*, 1–18. <https://doi.org/10.1080/00218499.2026.2620892>
 13. Dabiran, E., Farivar, S., Wang, F., & Grant, G. (2024). Virtually human: Anthropomorphism in virtual influencer marketing. *Journal of Retailing and Consumer Services*, 79, 103797.
<https://doi.org/10.1016/j.jretconser.2024.103797>
 14. De Veirman, M., Cauberghe, V., & Hudders, L. (2017). Marketing through Instagram influencers: The impact of number of followers and product divergence on brand attitude. *International Journal of Advertising*, 36(5), 798–828. <https://doi.org/10.1080/02650487.2017.1348035>
 15. Elyamany, N., Youssef, Y. O., & El-karef, N. (2025). A trans-disciplinary forensic study of Lil Miquela’s virtual identity performance in Instagram. *AI & Society*, 40(6), 4831–4854.
<https://doi.org/10.1007/s00146-025-02219-8>
 16. Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism.. *Psychological Review*, 114(4), 864–886. <https://doi.org/10.1037/0033-295x.114.4.864>
 17. Evans, N. J., Phua, J., Lim, J., & Jun, H. (2017). Disclosing Instagram influencer advertising: The effects of disclosure language on advertising recognition, attitudes, and behavioral intent. *Journal of Interactive Advertising*, 17(2), 138–149. <https://doi.org/10.1080/15252019.2017.1366885>
 18. Feng, Y., Chen, H., & Xie, Q. (2024). AI influencers in advertising: The role of AI influencer-related attributes in shaping consumer attitudes, consumer trust, and perceived influencer–product fit. *Journal of Interactive Advertising*, 24(1), 26–47. <https://doi.org/10.1080/15252019.2023.2284355>
 19. Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research*, 21(1), 1–31. <https://doi.org/10.1086/209380>

20. Google for Developers (2026). Commentthreads: List. *YouTube Data API*.
<https://developers.google.com/youtube/v3/docs/commentThreads/list>
21. Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267–297.
<https://doi.org/10.1093/pan/mps028>
22. Gutuleac, R., Baima, G., Rizzo, C., & Bresciani, S. (2024). Will virtual influencers overcome the uncanny valley? The moderating role of social cues. *Psychology & Marketing*, 41(7), 1419–1431.
<https://doi.org/10.1002/mar.21989>
23. Hannigan, T. R., Haans, R. F. J., Vakili, K., Tchalian, H., Glaser, V. L., Wang, M. S., Kaplan, S., & Jennings, P. D. (2019). Topic modeling in management research: Rendering new theory from textual data. *Academy of Management Annals*, 13(2), 586–632. <https://doi.org/10.5465/annals.2017.0099>
24. Horton, D., & Wohl, R. R. (1956). Mass communication and para-social interaction. *Psychiatry*, 19(3), 215–229. <https://doi.org/10.1080/00332747.1956.11023049>
25. Hussain, A., Narula, A., Singh, G., Sinha, M., Raj, R., & Nandal, N. (2026). Virtual influencers in digital marketing: A PRISMA-based systematic literature review and thematic synthesis of engagement, trust, and ethical challenges. *Future Business Journal*, 12(1), 210. <https://doi.org/10.1186/s43093-026-00910-w>
26. Karagür, Z., Becker, J., Klein, K., & Edeling, A. (2022). How, why, and when disclosure type matters for influencer marketing. *International Journal of Research in Marketing*, 39(2), 313–335.
<https://doi.org/10.1016/j.ijresmar.2021.09.006>
27. Khalfallah, D., & Keller, V. (2025). Authenticity, ethics, and transparency in virtual influencer marketing: A cross-cultural analysis of consumer trust and engagement: A systematic literature review. *Acta Psychologica*, 260, 105573. <https://doi.org/10.1016/j.actpsy.2025.105573>
28. Kim, M., & Baek, T. H. (2024). Are virtual influencers friends or foes? Uncovering the perceived creepiness and authenticity of virtual influencers in social media marketing in the United States. *International Journal of Human–Computer Interaction*, 40(18), 5042–5055.
<https://doi.org/10.1080/10447318.2023.2233125>
29. Kim, E., Kim, D., E, Z., & Shoenberger, H. (2023). The next hype in social media advertising: Examining virtual influencers' brand endorsement effectiveness. *Frontiers in Psychology*, 14, 1089051.
<https://doi.org/10.3389/fpsyg.2023.1089051>
30. Kinder-Kurlanda, K. E., & Weller, K. (2020). Research ethics practices in a changing social media landscape. *AoIR Selected Papers of Internet Research*. <https://doi.org/10.5210/spir.v2020i0.11251>
31. Koles, B., Audrezet, A., Moulard, J. G., Ameen, N., & McKenna, B. (2024). The authentic virtual influencer: Authenticity manifestations in the metaverse. *Journal of Business Research*, 170, 114325.
<https://doi.org/10.1016/j.jbusres.2023.114325>
32. Leung, F. F., Gu, F. F., Li, Y., Zhang, J. Z., & Palmatier, R. W. (2022). Influencer marketing effectiveness. *Journal of Marketing*, 86(6), 93–115. <https://doi.org/10.1177/00222429221102889>

33. Lim, R. E., & Lee, S. Y. (2023). "You are a virtual influencer!": Understanding the impact of origin disclosure and emotional narratives on parasocial relationships and virtual influencer credibility. *Computers in Human Behavior*, 148, 107897. <https://doi.org/10.1016/j.chb.2023.107897>
34. Liu, F., & Wang, R. (2025). Fostering parasocial relationships with virtual influencers in the uncanny valley: Anthropomorphism, autonomy, and a multigroup comparison. *Journal of Business Research*, 186, 115024. <https://doi.org/10.1016/j.jbusres.2024.115024>
35. Looi, J., Kim, E., & E, Z. (2025). Sponsorship disclosure in virtual influencer marketing: Assessing users' sentiment and engagement toward virtual influencer endorsements. *Journal of Advertising Research*, 65(3), 321–343. <https://doi.org/10.1080/00218499.2025.2464300>
36. Lou, C., & Yuan, S. (2019). Influencer marketing: How message value and credibility affect consumer trust of branded content on social media. *Journal of Interactive Advertising*, 19(1), 58–73. <https://doi.org/10.1080/15252019.2018.1533501>
37. Marani, A. H., & Baumer, E. P. S. (2024). A review of stability in topic modeling: Metrics for assessing and techniques for improving stability. *ACM Computing Surveys*, 56(5), 1–32. <https://doi.org/10.1145/3623269>
38. Mirowska, A., & Arsenyan, J. (2023). Sweet escape: The role of empathy in social media engagement with human versus virtual influencers. *International Journal of Human-Computer Studies*, 174, 103008. <https://doi.org/10.1016/j.ijhcs.2023.103008>
39. Mori, M., MacDorman, K., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*, 19(2), 98–100. <https://doi.org/10.1109/mra.2012.2192811>
40. Mouritzen, S. L. T., Penttinen, V., & Pedersen, S. (2024). Virtual influencer marketing: The good, the bad and the unreal. *European Journal of Marketing*, 58(2), 410–440. <https://doi.org/10.1108/ejm-12-2022-0915>
41. Munnukka, J., Maity, D., Reinikainen, H., & Luoma-aho, V. (2019). 'Thanks for watching': The effectiveness of YouTube vlog endorsements. *Computers in Human Behavior*, 93, 226–234. <https://doi.org/10.1016/j.chb.2018.12.014>
42. Nguyen, B. P., & Wu, W. (2026). How and when do virtual influencers work? A meta-analysis of mechanisms and moderators in digital commerce. *Journal of Theoretical and Applied Electronic Commerce Research*, 21(4), 124. <https://doi.org/10.3390/jtaer21040124>
43. Ohanian, R. (1990). Construction and validation of a scale to measure celebrity endorsers' perceived expertise, trustworthiness, and attractiveness. *Journal of Advertising*, 19(3), 39–52. <https://doi.org/10.1080/00913367.1990.10673191>
44. Oliveira, A. B. d. S., & Chimenti, P. (2021). 'Humanized robots': A proposition of categories to understand virtual influencers. *Australasian Journal of Information Systems*, 25. <https://doi.org/10.3127/ajis.v25i0.3223>
45. Pan, M., Blut, M., Ghiassaleh, A., & Lee, Z. W. Y. (2025). Influencer marketing effectiveness: A meta-analytic review. *Journal of the Academy of Marketing Science*, 53(1), 52–78. <https://doi.org/10.1007/s11747-024-01052-7>

46. Pushparaj, P., Kushwaha, B. P., & Prashar, S. (2025). A systematic literature review of virtual influencers in marketing using bibliometric analysis. *International Review on Public and Nonprofit Marketing*, 22(3), 631–662. <https://doi.org/10.1007/s12208-025-00441-0>
47. Reinikainen, H., Munnukka, J., Maity, D., & Luoma-aho, V. (2020). ‘You really are a great big sister’ – parasocial relationships, credibility, and the moderating role of audience comments in influencer marketing. *Journal of Marketing Management*, 36(3–4), 279–298. <https://doi.org/10.1080/0267257x.2019.1708781>
48. Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B., & Rand, D. G. (2014). Structural topic models for open-ended survey responses. *American Journal of Political Science*, 58(4), 1064–1082. <https://doi.org/10.1111/ajps.12103>
49. Roberts, M. E., Stewart, B. M., & Tingley, D. (2019). STM: An R package for structural topic models. *Journal of Statistical Software*, 91(2). <https://doi.org/10.18637/jss.v091.i02>
50. Rohde, P., & Mau, G. (2021). “It’s selling like hotcakes”: Deconstructing social media influencer marketing in long-form video content on YouTube via social influence heuristics. *European Journal of Marketing*, 55(10), 2700–2734. <https://doi.org/10.1108/ejm-06-2019-0530>
51. Sands, S., Campbell, C. L., Plangger, K., & Ferraro, C. (2022). Unreal influence: Leveraging AI in influencer marketing. *European Journal of Marketing*, 56(6), 1721–1747. <https://doi.org/10.1108/ejm-12-2019-0949>
52. Schwemmer, C., & Ziewiecki, S. (2018). Social media sellout: The increasing role of product promotion on YouTube. *Social Media + Society*, 4(3), 2056305118786720. <https://doi.org/10.1177/2056305118786720>
53. Sesar, V., Martinčević, I., & Boguszewicz-Kreft, M. (2022). Relationship between advertising disclosure, influencer credibility and purchase intention. *Journal of Risk and Financial Management*, 15(7), 276. <https://doi.org/10.3390/jrfm15070276>
54. Song, X., Lu, Y., & Yang, Q. (2024). The negative effect of virtual endorsers on brand authenticity and potential remedies. *Journal of Business Research*, 185, 114898. <https://doi.org/10.1016/j.jbusres.2024.114898>
55. Sorosrungruang, T., Ameen, N., & Hackley, C. (2024). How real is real enough? Unveiling the diverse power of generative AI-enabled virtual influencers and the dynamics of human responses. *Psychology & Marketing*, 41(12), 3124–3143. <https://doi.org/10.1002/mar.22105>
56. Stein, J., Breves, P. L., & Anders, N. (2024). Parasocial interactions with real and virtual influencers: The role of perceived similarity and human-likeness. *New Media & Society*, 26(6), 3433–3453. <https://doi.org/10.1177/14614448221102900>
57. Tian, M., Hu, H., & Chen, M. (2026). Virtual influencers in interactive marketing: A state-of-art review and future research directions. *Journal of Research in Interactive Marketing*, 20(6), 794–811. <https://doi.org/10.1108/jrim-03-2025-0123>
58. Zhou, B., Li, J., Abdullah, Z., Hashim, N., & Qi, M. (2026). Virtual influencer attributes and purchase intention: A systematic review. *Frontiers in Human Dynamics*, 8, 1816209.

59. Zniva, R., Weitzl, W. J., & Lindmoser, C. (2023). Be constantly different! How to manage influencer authenticity. *Electronic Commerce Research*, 23(3), 1485–1514. <https://doi.org/10.1007/s10660-022-09653-6>

Figures




Figure 1

Corpus construction and primary STM sample. The figure shows the progression from qualifying official-channel and brand-uploaded videos to the final English-language, top-level-comment corpus. Replies were retained for sensitivity analysis; later duplicate text and comments containing no retained terms were removed




Figure 2

Monthly qualifying comment volume and sponsorship support. Counts use the 2,014 English-language top-level comments satisfying the primary eligibility rules. Sponsored-video comments first occur in 2023, precluding a supported full-window time-by-sponsorship interaction




Figure 3

Topic-count diagnostic comparison. Candidate models with K from 8 to 25 are compared using semantic coherence, exclusivity, held-out likelihood and residual dispersion. K = 15 was selected from joint diagnostics and representative-document interpretability




Figure 4

Estimated topic-correlation network. Node area represents mean topic prevalence; plotted edges are positive conditional correlations of at least .05; colours show Louvain communities. The network is descriptive, not causal




Figure 5

Model-estimated prevalence of research-facing topics over time. Curves show adjusted expected topic proportions; shaded areas are 95% simulation intervals. Estimates control for sponsorship, comment likes and comment delay




Figure 6

Qualifying upload events across the study window. The clustering of official uploads in 2021, sponsored activity in 2023–2024 and renewed official uploads in 2026 show why trends require interpretation alongside content events




Figure 7

Adjusted sponsorship associations across all topics. Points show sponsored-minus-non-sponsored expected-prevalence differences; horizontal lines are 95% simulation intervals. Differences are associational




Figure 8

Specification sensitivity of sponsorship contrasts. Selected contrasts are compared across primary and alternative specifications after one-to-one topic alignment. Weakly aligned solutions require additional caution