Journal of Cyberspace Studies
eISSN: 2588-5502
http://jcss.ut.ac.ir
University of Tehran

Little digital natives and the public moral order of Persian social media

Kosar Saadatnia[1]*, Bijan Tafazzoli[2]

1. Department of Communication Sciences, Faculty of Social Sciences, University of Tehran, Tehran, Iran. (*Corresponding author: ✉ kosar.saadatniya@gmail.com, https://orcid.org/0009-0000-3382-6261).
2. Department of International Relations, Graduate School of Social Sciences and Humanities, Koç University, Istanbul, Turkiye.

Article Info

Original article

Main Object:
Humanities & Social Sciences

Received: 03 July 2026
Revised: 17 July 2026
Accepted: 18 July 2026
Published online: 07 September 2026

Keywords:
child influencers,
children and the internet,
computational communication,
digital natives,
framing,
gender,
lateral family surveillance,
Persian social media,
topic modeling.

Abstract

Background: Although the "digital native" has been empirically discredited as a scientific category, it survives as a folk figure through which adults register wonder and alarm at children's platform lives. Research on children and the internet has relied largely on Western, survey-based studies of children's experiences and parents' reported concerns, or on episodic moral panics amplified by professional media.

Aims: This study maps how Persian-language Twitter users publicly construct the intersection of children and the internet, examining the temporal architecture of the discourse across the September 2022 blocking of Instagram in Iran, its thematic structure, its frames of concern and their gendering, and the engagement dynamics that determine which talk travels.

Methodology: A corpus of 1,467 Persian tweets posted by 1,385 accounts between February 2021 and May 2025 was analyzed by combining non-negative matrix factorization topic modeling (benchmarked against LDA and a co-word network decomposition), dictionary-based frame analysis, PELT change-point detection, and negative binomial engagement regressions, with the full corpus closely read and all excerpts translated by the author.

Findings: The dominant genre is lateral family surveillance: young adults publicly exhibiting the online conduct of younger relatives, with an exact child age stated in 31.6% of all posts. The 2022 blocking contracted the discourse by 38% before it recovered above baseline; the children-in-protest frame briefly saturated the corpus and then evaporated. Concern is sharply gendered, 96% of sexualization posts and 88% of child-influencer posts are girl-marked, the influencer economy is the only growing frame, and state child-protection vocabulary appears in 0.5% of posts. Kinship anecdotes drive conversational reciprocity; policy indignation does not.

Conclusion: The public moral order of Persian social media is constructed from the family outward, not from the state downward.

Cite this article: Saadatnia K, Tafazzoli B. (????). "Little digital natives and the public moral order of Persian social media". *Cyberspace Studies*. ?(?): 1-20. doi: https://doi.org/10.22059/jcss.2026.417807.1259.

Cyberspace Studies, Vol ?, No ?, ????

1. Introduction

Few phrases have organized public feeling about childhood and technology as durably as the "digital native". Coined to describe a generation of learners presumed to be constitutionally different because they were born into networked media (Prensky, 2001), the concept was dismantled by empirical research within a decade: age, it turned out, predicted digital competence far less well than experience, breadth of use, and education (Helsper & Eynon, 2013). Yet the figure survived its refutation. It lives on not as a scientific category but as a folk one; a shorthand through which adults everywhere register astonishment, amusement, and alarm at the sight of a nine-year-old navigating a smartphone with unsettling fluency. This paper is about that folk category at work. It examines how a large non-Western public talks, in its own language and on its own platforms, about its little digital natives: the toddler whose WhatsApp status is a K-pop idol, the eleven-year-old whose Instagram story becomes family news, the child influencer with forty thousand followers before her third birthday. What such talk produces, we argue, is a public moral order; a dispersed, cumulative, publicly performed set of judgments about what is normal, laughable, worrying, or shameful in children's platform lives, assembled one post at a time rather than decreed by any authority.

Two scholarly traditions have approached adult concern about children online, and the present study is positioned in the gap between them. The first is the risk-and-opportunity tradition exemplified by the EU Kids Online network, which has documented, mostly through cross-national surveys of children and parents, how online risks are distributed, how they relate to opportunities, and how parental mediation and digital skills shape outcomes (Livingstone & Helsper, 2013; Livingstone et al., 2018; Livingstone & Stoilova, 2021). The second is the moral panic tradition, which treats adult anxiety about children and new media as episodic eruptions, "technopanics", amplified by professional media and translated into restrictive legislation (Cohen, 2011; Marwick, 2008). Both traditions are indispensable, and both share a blind spot. The first elicits concern through research instruments, asking parents and children to respond to categories the researcher supplies (Greyson et al., 2023); the second studies concern at its media-fueled peaks. Neither observes the sustained, low-intensity, bottom-up talk through which ordinary people constitute the connected child as a public object in the flow of everyday life; and almost nothing of this kind has been attempted outside Western, English-speaking settings. What a vernacular public actually says about children and the internet when nobody has asked it anything remains, with few exceptions, unmapped.

Iran offers a case of unusual analytical value for such a mapping. It combines an intensely connected society with one of the world's most restrictive internet governance regimes, in which everyday platform

access is mediated by filtering, throttling, and mass circumvention, and in which using a VPN has become an ordinary competence rather than an activist skill (Dal & Nisbet, 2022). Instagram, with an estimated forty million-plus Iranian users before its blocking, functioned as the country's de facto infrastructure for commerce, celebrity, and family sociality; children's internet use, meanwhile, has been a headline object of elite policy, most visibly in the parliamentary "Protection Bill" that proposed to restructure the national internet in the name of shielding minors. On 21 September 2022, following the nationwide protests triggered by the death of Mahsa (Zhina) Amini in morality police custody, Instagram and WhatsApp were blocked; a sharp, precisely dated rupture in the infrastructure of everyday discourse. Children and adolescents were central figures of that uprising, as schoolgirl protesters and as casualties whose names circulated globally. The period studied here thus contains, in a single arc, the ordinary child of the family feed, the extraordinary child of the revolution, and a state intervention justified in the child's name.

Framing theory supplies the conceptual apparatus. Frames select and make salient particular aspects of reality so as to promote a problem definition, a causal diagnosis, a moral evaluation, and a remedy (Entman, 1993), and social media research has shown that lay users are active framers rather than passive receivers of elite definitions (Mendelsohn et al., 2021). Computational methods now make such lay framing tractable at corpus scale (van Atteveldt & Peng, 2018). The analyses reported here combine non-negative matrix factorization (Lee & Seung, 1999), benchmarked against latent Dirichlet allocation (Blei et al., 2003), with a theory-driven Persian lexical dictionary of concern frames and negative binomial models of engagement. Because pretrained Persian-language classifiers of adequate quality remain scarce, algorithmic outputs were treated throughout as hypotheses to be checked against the texts: the full corpus was closely read, and all quoted material appears in the author's English translation. The combination trades scalability for interpretive depth, which is, for a vernacular corpus of this size and language, the defensible trade.

Two further bodies of scholarship converge on the corpus and sharpen its significance. Research on sharenting and the child-influencer economy has documented how children's images have become raw material for parental self-presentation and, increasingly, for monetized content production, raising concerns about labor, consent, and privacy (Blum-Ross & Livingstone, 2017; Abidin, 2015; Feller & Burroughs, 2021; Divon et al., 2026). That literature has been built almost entirely on content and regulation in the Global North; this study observes the same economy from the side of its vernacular reception in Iran, where the monetized child turns out to be the only frame of concern that grows across the study period. Meanwhile, the computational literature on Persian social media has concentrated on

Cyberspace Studies, Vol ?, No ?, ????

elections, protest, and crisis communication (Kermani & Tafreshi, 2022; Kermani, 2023; Shokrollahi et al., 2021). By taking up an everyday, ostensibly non-political topic and following it through a major political rupture, the present study asks what such episodes do to the durable structures of ordinary public talk; and finds that they interrupt without transforming them.

Against this background, the present study analyzes a corpus of Persian-language tweets posted by Iranian accounts, asking how Persian-speaking Twitter users publicly construct the intersection of children and the internet. The unit of interest is deliberately not the child's own experience, nor the parent's privately reported worry, but the public act of commentary: who watches which children, what behaviors are singled out for exhibition or censure, which moral vocabularies are reached for, and what kinds of talk the platform's attention economy rewards. In this sense, the study belongs to the analysis of vernacular expertise; the informal but patterned way in which a lay public defines a social problem and polices its boundaries.

2. Methodology

The corpus was assembled through keyword-based retrieval of public Persian-language posts on Twitter/X, using queries that combined common Persian terms for children (بچه, دختر بچه, پسر بچه and related kinship-age expressions) with the names of major platforms (اینستا/اینستاگرام, توییتر, تلگرام, واتساپ). After deduplication, the final dataset comprised 1,467 unique tweets posted by 1,385 distinct accounts between 19 February 2021 and 15 May 2025, each retaining its full text, timestamp, and engagement metadata (like, retweet, and reply counts); view counts were discarded because the capture procedure truncated them at 990, below the observed maximum of 897 likes, rendering them unusable as an outcome variable. Texts were preprocessed by removing URLs, emoji, and Latin-script intrusions and by applying a custom Persian stopword list, producing a cleaned field for modeling alongside the untouched original. Because keyword collection inevitably admits noise, a manual relevance audit was performed on a random subsample (n= 40), which classified 55% of items as squarely on topic, 15% as tangential (child terms used metaphorically, or adults discussed in child-related framings), and 30% as off-topic keyword artifacts; all prevalence figures reported below are therefore conservative lower bounds, and the dictionary-based estimates, which condition on substantive vocabulary, are more robust to this dilution than raw volume counts. Ethical handling followed established practice for public social media text: no attempt was made to identify users, quoted posts appear only in the authors' English

translation with personal names and handles removed or generalized, and renderings stay close to the vernacular register of the originals, with profanity lightly attenuated.

Analysis proceeded in four linked stages. Temporal structure was examined by aggregating posts into a monthly series and applying PELT change-point detection to the variance-stabilized (square-root-transformed) counts, with the penalty parameter swept across values of 3, 5, and 8 to establish the robustness of detected breaks. Thematic structure was mapped with non-negative matrix factorization (Lee & Seung, 1999) fitted to a TF-IDF representation of the cleaned texts (minimum document frequency of 6, maximum of 40%, sublinear term-frequency weighting) across solutions of k= 4 to 12, with normalized pointwise mutual information coherence guiding model selection; the retained nine-topic solution was triangulated against a latent Dirichlet allocation benchmark (Blei et al., 2003) and a Louvain community decomposition of the PMI-weighted word co-occurrence network. Frames of concern were operationalized through a seven-category Persian lexical dictionary— kinship anecdote, parenting and control, child-influencer economy, sexualization and safety, children in protest, hijab on children, and filtering and platform policy— developed iteratively through close reading and regular-expression refinement, and coded as binary, non-mutually-exclusive indicators per tweet, with gender marking measured by the presence of the tokens دختر (girl) and پسر (boy); exact child ages were extracted by matching Persian and Latin numerals as well as spelled-out number words adjacent to year-of-age expressions, retaining values between 2 and 17. Finally, engagement was modeled with negative binomial (NB2) regressions estimated by maximum likelihood with heteroskedasticity-robust (HC1) standard errors, fitted separately to like and reply counts, with the seven frames, an exact-age indicator, question-mark presence, hashtag and mention counts, log-transformed text length, and year fixed effects (reference: 2021) as predictors; likelihood-ratio tests decisively rejected Poisson alternatives in both models (α= 3.73 for likes, 4.38 for replies, both $P< 0.001$). Throughout, algorithmic outputs were treated as hypotheses rather than conclusions: the full corpus was closely read, topic and frame assignments were verified against the texts, and the quantitative and interpretive strands were required to converge before any pattern was reported.

3. Findings

The final corpus comprised 1,467 Persian-language tweets posted by 1,385 unique accounts between 19 February 2021 and 15 May 2025. Instagram was the referential center of gravity of the discourse: 67.3% of tweets named the platform explicitly, against 6.0% for Twitter itself, 2.5% for Telegram, and 1.7% for WhatsApp, confirming that when

Cyberspace Studies, Vol ?, No ?, ????

Persian-speaking users talk about children online, they are overwhelmingly talking about children on Instagram. Because the corpus was assembled through keyword matching, a manual relevance audit was conducted on a random subsample (n= 40): 55% of items were squarely on topic, a further 15% were tangential (child terms used metaphorically, or adults discussed in child-related framings), and 30% were off-topic keyword artifacts. All proportions reported below should therefore be read as conservative lower bounds on the topical population, and the dictionary-based estimates, which select on substantive vocabulary, are more robust to this dilution than raw volume counts. View counts were excluded from all analyses because the capture procedure truncated them (maximum recorded value of 990 against a maximum of 897 likes), rendering them unusable as an outcome variable. Throughout the chapter, sample posts are presented in the authors' English translation; renderings stay close to the vernacular register of the originals, with profanity lightly attenuated, and identifying details (personal names and handles) removed or generalized.

3.1. Corpus overview and the shape of attention

Participation in the discourse was radically dispersed rather than organized: 95.9% of accounts appear exactly once in the corpus, and the most active account contributed only ten tweets. This is a crowd-sourced, episodic conversation, not the product of an advocacy field or a stable community of commentators. Engagement, by contrast, was extremely concentrated (Table 1).

Table 1. Engagement descriptives for the corpus (N = 1,467 tweets, 1,385 accounts)

Metric	M	Mdn	Max	Zero (%)
Likes	32.9	4	897	21.1
Retweets	4.0	0	490	82.4
Replies	4.4	1	315	44.8
Text length (characters)	182.8	184	—	—

Note. Data span 19 February 2021 to 15 May 2025. The top decile of tweets accounts for 79.0% of all likes. View counts were excluded owing to truncation at capture.

The mean tweet attracted 32.9 likes but the median only 4, one in five tweets (21.1%) received no likes at all, and the top decile of tweets absorbed 79.0% of all likes in the corpus. Retweeting was rarer still (82.4% of tweets were never retweeted), while replying was the most evenly distributed behavior (median= 1), a first indication, developed below, that this discourse circulates more through conversation than through amplification. Hashtags appeared in only 7.6% of tweets and user mentions in 1.9%, consistent with a vernacular register rather than campaign communication; 18.2% of tweets posed a question.

The one-off character of participation is worth illustrating, because

it shapes the very epistemology of the corpus: most users surface exactly once, deposit a single observation, confession, or query drawn from their private social world, and disappear. One such user wrote: "Years ago I followed a mother on Instagram who lived in Holland and had two children…, and her daughter suffered from a grave illness. I truly want to know how they are; if anyone knows her, send me the page. It was not a small page either. I do not know why, but for some time now my heart has been with that little girl" (October 2021, 3 likes).

Nothing suggests the authors ever returned to the conversation. The tweet is a message in a bottle, and the corpus, in aggregate, is an ocean of such bottles: brief, situated, affectively charged reports from a thousand private observatories.

3.2. Temporal dynamics: A discourse interrupted

Figure 1 plots monthly tweet volume. Change-point detection (PELT on the variance-stabilized series) identified a structural break in September 2022 that was robust across penalty values, partitioning the series into three regimes. From February 2021 through September 2022, the discourse ran at a mean of 28.4 tweets per month. Between October 2022 and July 2023 it collapsed to 17.5 tweets per month, a 38% contraction that coincides precisely with the blocking of Instagram and WhatsApp in Iran on 21 September 2022 and with the nationwide protests following the death of Mahsa (Zhina) Amini. From August 2023 onward the conversation not only recovered but exceeded its pre-block baseline (32.9 tweets per month, +16%), peaking at 48 tweets in both July and December 2024, a pattern consistent with the routinization of VPN-mediated access and the re-normalization of Instagram in everyday Iranian life despite formal filtering.

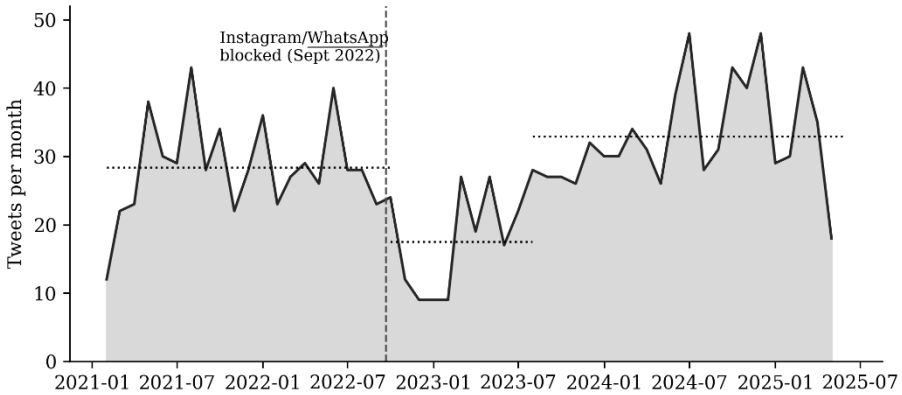

Note. Dotted horizontal lines mark segment means estimated by PELT change-point detection (28.4, 17.5, and 32.9 tweets per month). The dashed vertical line marks the blocking of Instagram and WhatsApp in Iran (21 September 2022).

Figure 1. Monthly tweet volume, February 2021 to May 2025

In the pre-block regime the conversation runs in its default comic-domestic key, as in this February 2021 report: "My 10-year-old cousin has posted a story saying GTA taught him that 'life is a film and we are its actors.' Honestly, the only things I ever learned from GTA were swearing and picking up girls."

The trough that follows was not merely quantitative but qualitative. Within the protest window of 16 September to 31 December 2022, 54.2% of all corpus tweets carried the children-in-protest frame, and this window alone contains 32 of the 62 protest-frame tweets in the entire corpus. The affective temperature of those weeks is captured by a user who cut ties over an ill-timed story: "Today I blocked a close friend of several years because he posted a story of his café outing. You would have to be an utter ingrate, utterly stone-hearted, to be thinking of showing off your coffee while they are killing 15- and 16-year-old girls in the street" (October 2022).

Children entered the discourse as protagonists and casualties of the uprising, and celebrity Instagram became a relay of their suffering and defiance: "A famous former footballer has shared a story: a little girl on the Chalus road, under a tent, shivering, her voice trembling too—'we are frozen under this tent'—bruises under her eyes from cold, hunger, and every kind of deficiency; and I felt my heart catch fire. #MahsaAmini" (January 2023).

Another user relayed, approvingly, that "a well-known actress has posted an interview with a little girl on Instagram—it is very good. #MahsaAmini #NikaShakarami." Even the child's own artifacts circulated as protest matériel: "I saw on Instagram a little American girl's drawing of Zhina … there is so much said inside that frame."

By the recovery regime the default key returns, now audibly inflected by comparison and envy: "Yeeees—here is the story of my 17-year-old cousin, and here is me, who must study five more years before I can even drag the title 'doctor' behind my name. I do not know whether to thank my parents that I did not turn out like them, or to grieve that I grew up under restriction …" (July 2024).

The protest frame, meanwhile, evaporates almost completely, falling to 2.3% of tweets in 2024 and 0.6% in 2025 (Figure 2), a striking illustration of how quickly the extraordinary child of the uprising was displaced by the ordinary child of the feed.

3.3. Thematic structure

To map the latent organization of the discourse, non-negative matrix factorization (NMF) was fitted to a TF-IDF representation of the preprocessed texts across solutions of 4 to 12 topics, with normalized pointwise mutual information (NPMI) coherence guiding model selection.

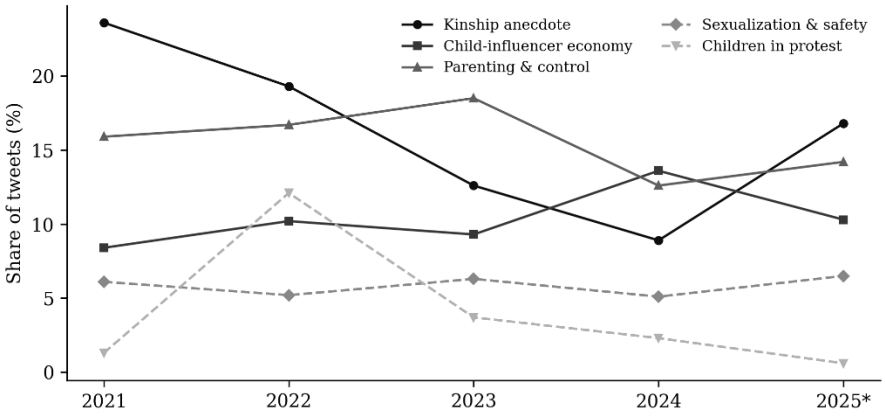

Note. Shares are percentages of all tweets in each calendar year. 2025 covers January to mid-May only.

Figure 2. Yearly prevalence of concern frames, 2021–2025

The nine-topic solution was retained as offering the best balance of coherence (NPMI= −0.13, within 0.01 of the maximum) and interpretability; an LDA benchmark and a Louvain community decomposition of the PMI-weighted word co-occurrence network recovered substantively the same structure (kinship storytelling, platform ecology, parental control, profile-picture nostalgia, and a distinct Mahsa Amini–Nika Shakarami cluster), lending triangulated support to the partition. Table 2 presents the nine topics.

The single most consequential feature of this solution is that four of the nine topics (Topics 3, 5, 8, and 9, jointly 29.0% of the corpus) are variants of one recognizable genre: adults reporting, mostly with amusement or mock despair, on the online conduct of the children in their extended family. A fifth topic (Topic 6, 9.0%) inverts the same gaze onto the speaker's own adolescent past. Only one topic (Topic 7, 8.9%) is political in the conventional sense. The everyday discourse of "internet and children" in this corpus is thus dominated not by policy, panic, or pedagogy, but by the family watching itself.

A brief tour of sample posts conveys the texture behind the labels. Topic 1, the broad everyday-feed topic, is where children appear as ambient content and as objects of longing: "I saw one reel on Instagram, and right now I would give everything I have to have a little girl and dote on her" (July 2024).

Another user reframes the same consumption as emancipation: "It is good that at this age, instead of raising a child, I watch lovely clips of children on Instagram and delight in them. That is what I call progress for the Iranian girl, against the talk of people who, rather than praying that she finds her footing and reaches her goals and passions, wish her a husband and a baby."

Table 2. NMF topic solution (k= 9): Labels, prevalence, and highest-loading terms

Topic	Label	Share (%)	Highest-loading terms (gloss, original)
1	Children in the everyday feed	30.2	page (پیج), clip (کلیپ), life (زندگی), mother (مادر), marriage (ازدواج)
2	Generational turf talk on Twitter	11.7	the kids of (بچه‌های), Twitter (توییتر), decade cohort (دهه), follow (فالو), account (اکانت)
3	Cousins' stories reported	11.4	boy (پسر), maternal uncle (دایی), paternal aunt (عمه), his/her stories (استوری‌های), relationship (رابطه)
4	Platform banter about "the kids"	11.2	kids (بچه‌ها), Twitter (توییتر), in love (عاشق), the space (فضای), account (اکانت)
5	Relatives' stories across apps	10.1	story (استوری), sister (خواهر), uncle's (عموی), WhatsApp (واتساپ), congratulations (تبریک)
6	Adolescent digital selves remembered	9.0	profile (پروفایل), photo (عکس), my Telegram (تلگرامم), I used to (می‌کردم), change (عوض)
7	Children of the 2022 uprising	8.9	post (پست), Mahsa (مهسا), Amini (امینی), Nika (نیکا), like (لایک)
8	Sibling onboarding and monitoring	4.2	little brother (داداش), installed (نصب), thirteen (سیزده), my followers (فالوم), wrote (نوشته)
9	Extended-family story watching	3.3	maternal aunt (خاله), story (استوری), husband (شوهر), age (سن), follow (فالو)

Note. Share= percentage of tweets for which the topic carries the largest document loading. Glosses are the authors' translations. Model selection by NPMI coherence over k= 4–12.

Topic 2 is boundary work against juvenile newcomers to Twitter: "Schools have reopened—what is this child doing here? Girl, go write your homework; this is not Instagram"; or, trailing off in disgust, "Everyone is busy with a handful of little girls who these days are turning Twitter, too, like Instagram, into… ."

Topic 3 is the cousin report proper, complete with folk hermeneutics: "My 12-, 13-year-old cousin posted a story: 'Life has taught me that nothing is unexpected from anyone.' Probably one of his teachers marked him down; there can be no other explanation for this story."

Topic 4 hosts vernacular platform theory of remarkable confidence: "On YouTube it is mostly the very young who comment; Instagram—

forget it; Twitter is mostly resentful kids whose age has crept up and who, having sprouted beards and breasts, believe they have grown up—the 'open-minded' 2000s-born, belching pseudo-intellectual nonsense."

Topic 5 registers the child's presence across the entire app ecology, down to the toddler: "My three-year-old sister's WhatsApp status is Jimin (:".

Topic 6 turns the same gaze backward onto the speaker's own adolescent digital self: "In this profile picture I was 16; how quickly three years have passed, man"; "I wish I were 15; I would put up a black-and-white profile photo with an aphorism, if only to lighten the load of my whining" (July 2024).

Topic 7 gathers the uprising material examined in the previous section.

Topic 8 documents onboarding and its informal gatekeeping: "Honestly, I was about to install Instagram for my thirteen-year-old brother, since so many of his friends are on it and he feels, in his own mind, that he is competing with them; today I learned I must wait another eight years" (March 2021).

And, Topic 9 completes the family panopticon from the maternal-aunt side of the ledger: "My 12-year-old cousin's profile picture vs. my profile picture" (August 2021, 48 likes); a one-line template in which the entire genre is compressed.

3.4. Watching, age-tagging, telling: Lateral family surveillance

The signature discursive practice of the corpus is what may be called lateral family surveillance: the public narration, by young adults, of what their younger siblings, cousins, nieces, and nephews are doing online. By the strict kinship dictionary this genre marks 15.7% of tweets (n= 230), and by topic assignment roughly 29%. Its canonical form is a highly compressed, almost formulaic report, "my 11-year-old cousin's story", sometimes with no comment at all beyond a screenshot, functioning as a recognizable meme template. Fuller instances stage the collision of generational registers: "My 9-year-old brother installed Instagram and sent me a follow request first thing in the morning. I don't know how one is supposed to react"; or, "My 13-year-old cousin posted a story: 'If you stand alone, you stand above kings.' Come down from there, child." What unites these reports is that the child does not speak; the child is exhibited.

The genre's registers run from deadpan exhibition through affectionate satire to something close to elegy. Satire elaborates the formula: "Since yesterday my 15-year-old cousin has posted forty birthday greetings from his friends as stories—every one of them 'sultan,' 'man of honor,' 'may you live a thousand years,' 'enjoy your love' :))))" (February 2021).

Self-mockery folds the watcher into the joke: "My 13-year-old cousin has posted a story—mahgardemoon mobarak ['happy

monthiversary to us']—while I am single. God, was this fair???" (April 2022, 119 likes, 22 replies).

The watcher can even be enrolled by the watched: "Once I saw a child's photo on Instagram, and my three-year-old cousin asked, who is this? I said, 'That is my daughter.' For three years now she has told everyone she meets that this one took a wife, has a child, and told no one."

The gaze also polices secrets laterally rather than vertically: "My 15-year-old cousin posts stories with his girlfriend and begs me not to tell the family! Me in that moment:" (March 2023).

And occasionally the deadpan cracks into something else entirely. The corpus's second most-liked tweet is the genre at its barest and most desolate—"The story of my 9-year-old cousin, who has just lost her mother" (March 2022, 851 likes); while another user reads a brother's caption as a family diagnosis: "The text of my 15-year-old brother's story: as a family, we grow up too fast, and that is not a good thing."

A striking feature of this practice is compulsive age accounting. Across the whole corpus, 31.6% of tweets state an exact child age between 2 and 17 (median= 11, interquartile range 7–14, with modal values at 15, 12, and 10 years), and within the kinship genre the figure rises to 77.0%. The stated age functions rhetorically as the point of the tweet: it calibrates exactly how scandalous, precocious, or endearing the reported behavior should be judged to be. Persian Twitter, in other words, does not simply say that children are online; it timestamps them, one birthday at a time. A complementary nostalgic genre performs the same accounting on the self: "I wish I were 13 again; I would turn my Telegram profile picture black to show the depth of my sorrow."

Here the adult speaker retrieves his or her own adolescent digital self as an object of affectionate ridicule, confirming that the watched child and the remembering adult are two moments of a single reflexive practice.

3.5. Frames of concern

Dictionary-based frame analysis (Table 3) locates the explicitly normative registers of the discourse. Parenting and control is the largest frame of concern (15.4%): tweets asking where the parents are, debating supervision and permission, and returning insistently to the child's phone as the object of contention. A December 2024 tweet (48 likes) states the position in full: "Amid these latest Twitter affairs—I mean the flasher pedophiles—my worry is this: how is it that the parents of a fourteen-year-old girl have no control whatsoever over the child's phone use? Why should a fourteen-year-old have a Twitter or Instagram account at all?"

Table 3. Prevalence and gender marking of concern frames (dictionary-based)

Frame	n	%	Girl-marked (%)	Boy-marked (%)
Kinship anecdote	230	15.7	66	48
Parenting and control	226	15.4	78	18
Child-influencer economy	156	10.6	88	12
Sexualization and safety	84	5.7	96	21
Children in protest	62	4.2	61	8
Hijab on children	32	2.2	94	19
Filtering and platform policy	7	0.5	—	—

Note. Frames are not mutually exclusive. Girl-marked and boy-marked = percentage of frame tweets containing the tokens دختر (girl) or پســـر (boy), respectively. Gender shares are omitted for the filtering frame owing to its very small n.

Blame is directed less at platforms than at parental strangeness: "The issue here is what harmful traces a child carries away from watching [a pop singer's] music video—while on Instagram a mother has tied a white headscarf on her four-year-old daughter and the two of them sing along to 'Somayeh, are you smoking weed?' The problem, graver than the internet or unsuitable singers, is that the parents are strange."

Parental restraint, conversely, is itself performed for the timeline: "She wanted to make a clip from the photos of the two girls and post it as her Instagram story and WhatsApp status. I did not allow it. Two of her closest friends cannot have children; it would not be right. #DaughtersDay" (June 2021, 50 likes).

The child-influencer economy follows at 10.6%, and it is the only frame that clearly grows over the period, reaching 13.6% of tweets in 2024 (Figure 2). Its tone is set by the very first tweet in the corpus: "Setting a little girl beside you in your videos and saying whatever you please to attract followers is a kind of fraud, done purely for money."

The register alternates between amusement and alarm. Amusement: "My cousin's little girl is two years and seven months old, and they have made her an Instagram account. She already has 40K followers. Sad 5 follower noises" (February 2021); "I follow so many little girls' pages on Instagram that I have no grown-ups' pages left." Alarm arrives when the observer is herself a mother: "There is a blogger on Instagram who has become famous through her color-eyed little daughter. She is my daughter's age—two. When I compare them, I ache for that child: if my girl spends two evenings running as a guest and her bedtime shifts, she turns cranky and frazzled—while that child is out with her mama every night until near dawn, purely for advertising" (July 2023).

At the frame's darkest edge, the monetized child fuses with child marriage: "Five or six pages in this little girl's name on Instagram—do not doubt it, the pages are not in her own hands; the real aim is attracting advertisement, and for visibility what sells better than child marriage?

Please repost, so that these damned accounts are taken down" (December 2024).

The longitudinal arc of the whole economy is condensed in a single confession: "Years ago on Instagram a page used to post the dances of a four- or five-year-old girl. I was young and foolish; I enjoyed it and followed. Then the page moved into advertising, and I slowly understood that this was the exploitation of a child, and I unfollowed. Today I came across it by chance: the girl is an adolescent now, dancing in a somewhat sexualized way, and her followers are very many" (March 2021).

Sexualization and safety, at 5.7%, is smaller but by far the most morally charged register; the single most-liked tweet in the corpus (897 likes) belongs to it: "You know what is interesting? When people watch that Instagram video where the father—apparently the stepfather—fondles the little girl's behind while hugging her and kisses her on the lips, they say: whyyy do you slap on the pedophile label, it is out of affection :)))" (February 2024).

The frame's argumentative core is a defense of categorical lines: "Nowhere in the world do they put a six-year-old child among four 'sexy' girls to perform suggestive moves on camera; nowhere do they tell a child, first Instagram, then books. Do you find it right that my six-year-old niece comes to me saying, 'come on, shake it, sexy'? There is a difference between a model and porn" (March 2021).

Threat and the ordinary run concurrently in the same timelines: "In the very midst of reading the stormy, nauseating, terrifying affair of the pedophiles on Twitter, two of my friends' little girls have opened Instagram accounts and invited me to their follower lists. Freedom has meaning only in the shelter of law" (June 2024)—one of the corpus's rare vernacular gestures toward regulation.

Occasionally the discourse of danger is even turned upon the self, the adult publicly documenting his own boundary-keeping: "A 16-year-old girl has latched onto me, a 34-year-old, on Instagram. How exactly am I supposed to handle this? I truly do not know, and I need help. #serious" (February 2021, 7 replies).

Hijab on children (2.2%) surfaces episodically around images of veiled small girls, and it is argued in both directions. Against: "A well-known clerical Instagram figure had posted a story of his little daughter—a two-year-old peanut with a headscarf on her head. Furious, I messaged him: what is this rag you have put on the child's head? He replied: she was simply cold. I was ashamed; though I do hope that when she is grown he will not compel her in anything. #CompulsoryHijab" (February 2021).

And for: "In Iraq they are very strict about hijab, and I respect them for it; their little girls would admonish me to put on the face veil, and I admire them. But then, in return, their people are not under cultural attack the way ours are—they do not have Instagram and the internet

and satellite at the intensity that exists here" (March 2021); a tweet in which, notably, the internet itself is cast as the invading force and the veiled child as the rampart.

Equally telling is what is nearly absent. Filtering and platform-policy vocabulary, including any reference to the much-debated "Protection Bill" (طرح صیانت), appears in only 7 tweets, 0.5% of the corpus, across more than four years; and close reading shrinks even this: two of the seven use "filter" in the photographic sense (a flowery Instagram effect, a Snapchat lens), leaving five genuinely policy-related posts. Their register, moreover, is fatalistic rather than deliberative. One user, reacting to a widely discussed tragedy involving a teenage girl, indicts every institution at once: "My heart aches for [her] father and mother; perhaps they wanted to raise a good daughter. Neither state television nor the schools are of any use, and in this wreck of a place Twitter is filtered while Instagram roams free, wretched the family that must raise a child in these conditions" (July 2021).

Another dismisses the policy instrument itself: "You short-sighted man—what effect does filtering even have? Everyone simply buys a VPN, and buys it easily; and once connected, alongside Instagram a thousand worse sites come within reach as well" (November 2024).

This, in a polity where internet governance for the sake of children has been a headline elite project, is the finding: the vernacular conversation about children and the internet and the official one run on separate tracks. Ordinary users moralize the family, the influencer, and the predator; almost never the regulator, and when they do, it is to declare the regulator beside the point.

3.6. The gendered child

The watched child of this discourse is overwhelmingly a girl. Tokens for "girl" occur in 73.6% of tweets against 21.1% for "boy", and although part of this base-rate asymmetry may reflect the composition of the collection query and should not be over-interpreted in absolute terms, the differential distribution of gender across frames within the corpus is robust and instructive. Girl references saturate precisely the protective and moralizing frames: 96% of sexualization-and-safety tweets, 94% of hijab tweets, 88% of child-influencer tweets, and 78% of parenting-and-control tweets are girl-marked, against boy shares of 21%, 19%, 12%, and 18%, respectively. The kinship-anecdote genre, by contrast, is the most gender-balanced register (66% girl-marked, 48% boy-marked), and boys appear there chiefly as comic protagonists, the swaggering 13-year-old with his aphoristic story captions. The pattern is unambiguous: girls enter the discourse as objects of endangerment, exposure, and honor, boys as subjects of amusement.

Users narrate this double standard explicitly, and from within. "To hell with a life in which my 15-, 16-year-old male cousin can go live on

Instagram drinking araq, while I, for fear of the relatives, dare not post one story out of line" (November 2021, 131 likes, 20 replies).

The policing extends backward into memory, in a retrospective that fuses the profile picture, the Snapchat lens, and the family telephone tree: "I was 15, ninth grade; my WhatsApp profile was a photo of me with my hair uncovered, taken with a Snapchat filter that made my lips a little fuller and red. The next evening my uncle telephoned my father: why has she put up such a photo—her lips red, her hair loose?" (May 2023, 115 likes).

It persists into adulthood: "God, if they find my account they will humiliate me before the whole family! They used to screenshot my Instagram stories and send them to my father—as if I were a little girl!" (April 2025).

And where boys' precocity is comic, girls' precocity is scandal: "Instagram is full of under-18 girls who have married, had children, opened a page, and quite cheerfully write up the story of how they met their gentleman. Seriously, what has happened to this 2000s-born generation? Damn it—you are not even permitted to open a bank account, and you go off and marry and bear children? God, put me through the grater" (May 2021).

The gendering of concern, in other words, reproduces offline moral economies of the female body on the terrain of the child's Instagram account.

3.7. Predictors of engagement

To test whether these frames govern the allocation of attention, negative binomial (NB2) regressions with heteroskedasticity-robust standard errors were fitted to like and reply counts, with the concern frames, age tagging, formal features, and year fixed effects as predictors (Table 4).

Overdispersion was severe ($\alpha = 3.73$ for likes and 4.38 for replies, both $P < 0.001$; likelihood-ratio tests decisively rejected Poisson), vindicating the specification. The likes model explains almost nothing (McFadden pseudo-$R^2 = 0.006$): with the exceptions of text length (IRR= 1.63 per log-character, $P < 0.001$) and the hijab frame, which depresses likes sharply (IRR= 0.32, $P = 0.009$), no content frame significantly predicts liking, and tweets from 2023 and 2024 attracted roughly 40% fewer likes than the 2021 baseline. Combined with the finding that the top decile of tweets absorbs 79% of all likes, this indicates a virality regime that is heavily concentrated yet essentially idiosyncratic: which child-related tweet explodes is not systematically foreseeable from what it says. The point is easily illustrated at the extremes: the fully argued position tweet on hijab, children, and cultural invasion quoted above earned zero likes, while the bare line about the newly bereaved nine-year-old earned 851.

Table 4. Negative Binomial (NB2) models of like and reply counts (n= 1,467)

Predictor	Likes, IRR [95% CI]	Replies, IRR [95% CI]
Kinship anecdote	1.00 [0.62, 1.63]	1.81 [0.94, 3.47]†
Child-influencer economy	0.81 [0.47, 1.37]	0.49 [0.30, 0.82]**
Parenting and control	0.88 [0.51, 1.53]	1.01 [0.60, 1.71]
Sexualization and safety	0.72 [0.30, 1.75]	0.53 [0.23, 1.21]
Children in protest	1.63 [0.75, 3.57]	1.12 [0.49, 2.55]
Hijab on children	0.32 [0.13, 0.75]**	2.13 [0.41, 10.97]
Exact age stated	0.94 [0.65, 1.35]	1.14 [0.75, 1.74]
Question mark	1.27 [0.86, 1.86]	0.88 [0.58, 1.33]
Hashtag count	0.95 [0.75, 1.21]	1.16 [0.82, 1.65]
Mention count	0.78 [0.55, 1.12]	0.75 [0.47, 1.17]
Log text length	1.63 [1.26, 2.12]***	1.41 [1.01, 1.98]*
Year 2022	0.94 [0.57, 1.54]	0.72 [0.42, 1.23]
Year 2023	0.59 [0.36, 0.96]*	0.56 [0.31, 1.00]†
Year 2024	0.63 [0.41, 0.98]*	0.94 [0.54, 1.65]
Year 2025	0.70 [0.43, 1.16]	0.74 [0.34, 1.64]

Note. IRR = incidence-rate ratio; CI = confidence interval based on robust (HC1) standard errors. Reference year is 2021. Dispersion α= 3.73 (likes) and 4.38 (replies), both $P< 0.001$; McFadden pseudo-R^2= 0.006 and 0.008. †$P< 0.10$. * $P< 0.05$. ** $P< 0.01$. *** $P< 0.001$.

Replies, the conversational face of engagement, are more patterned. Kinship anecdotes attract nearly twice the replies of other content (IRR= 1.81, $P= 0.076$; descriptively, their median of 7 likes also exceeds the 2–4 of every other frame), whereas tweets in the child-influencer-economy frame receive only half the replies of comparable tweets (IRR= 0.49, $P= 0.007$). The reply maximum of the entire corpus, 315 replies, belongs to a bare one-line prompt of exactly this kind, “My 12-year-old cousin’s profile picture:”, and the runners-up are of the same species: “My 15-year-old cousin’s story—and honestly, dirt on my head for being single” (October 2021, 741 likes, 114 replies); the bereaved-cousin tweet above (102 replies); and a father’s report (February 2022, 162 replies) that when a clip of dancing little girls surfaced on his sister’s Instagram feed, his own daughter, not yet three, announced “the most beautiful sentence of my life”: “I only dance at home, for my Papa”; for which, he adds, he pays out a heap of dance tips every night. The family anecdote, in short, is the conversational engine of this discourse, inviting recognition and reciprocal storytelling, while complaints about the monetized child circulate as broadcast indignation that others endorse but do not discuss. Stating the child’s exact age, despite its rhetorical centrality, confers no measurable engagement premium in either model.

Cyberspace Studies, Vol ?, No ?, ????

4. Conclusion

What, finally, do Persian-speaking users do when they talk about their little digital natives? Neither of the inherited scripts fits. The child of this corpus is not the precociously competent native of the generational imagination (Prensky, 2001), to be celebrated, nor the imperiled innocent of the technopanic tradition (Marwick, 2008), to be rescued by legislation. It is a domestic spectacle: a figure watched laterally by an extended family that narrates, age-stamps, and exhibits its conduct for an audience of strangers, and watched retrospectively by adults who turn the same amused gaze on their own adolescent digital selves. The public moral order of the title is therefore not a metaphor but an empirically observable division of moral labor, assembled one post at a time through micro-genres of everyday storytelling— the cousin's story, the age tag, the nostalgic profile picture— rather than through panic episodes or policy argument. Its mechanism is almost Durkheimian, producing solidarity through shared laughter and shared indignation, but its objects are anything but even-handed: the girl carries nearly the whole weight of protection, honor, and market anxiety, appearing in 96% of sexualization posts and 88% of child-influencer posts, while the boy furnishes comedy. And the platform's attention economy rewards precisely the genre that performs this regulation, directing conversational reciprocity toward the family anecdote and withholding it from broadcast complaint. In vernacular hands, the digital native turns out to be less a generational category than a moral occasion; a recurring pretext on which a society rehearses, enforces, and occasionally mocks its own norms about childhood, gender, and exposure.

The temporal and political findings sharpen rather than soften this claim. The autumn of 2022 shows the order's porousness and its resilience in a single arc: for a few months the watched cousin was displaced by the martyred schoolgirl, the blocking of Instagram cut away the discourse's infrastructural ground and volume contracted by 38%, and yet within a year the protest frame had dwindled to fractions of a percent while the everyday register returned above its former baseline. Extraordinary politics can commandeer vernacular attention; it does not rewrite vernacular grammar. The state, meanwhile, is the great absentee of its own project. A government that justifies its most ambitious internet legislation in the name of the child appears in the child's vernacular public in one-half of one percent of posts, and even then chiefly as an object of fatalism; an arrangement in which authoritarian internet governance commands infrastructure without commanding meaning, and in which policy aimed at children's online lives passes clean over the heads of the public that actually regulates them. These conclusions are bounded by the corpus that yields them: Iranian Twitter skews urban and oppositional, keyword collection admits noise, and a corpus of watchers can say nothing about the

experience of the watched child herself. Within those bounds, however, the study offers lateral family surveillance as a portable concept, and the obvious next question is comparative: whether extended-kinship societies elsewhere in the Middle East, South Asia, and beyond regulate their little digital natives through the same laughing, watchful, gendered domestic gaze, or whether the family panopticon documented here is a specifically Iranian formation shaped by a blocked internet and a moralizing state. Either answer would deliver something the survey tradition cannot: an account of how the moral life of connected childhood is actually lived; locally, publicly, and one story at a time.

Conflict of interest

The authors declared no conflicts of interest.

Ethical considerations

The authors have completely considered ethical issues, including informed consent, plagiarism, data fabrication, misconduct, and/or falsification, double publication and/or redundancy, submission, etc. This article was not authored by artificial intelligence.

Data availability

The dataset generated and analyzed during the current study is available from the author on reasonable request.

Funding

This research did not receive any grant from funding agencies in the public, commercial, or non-profit sectors.

References

Abidin, C. (2015). "Communicative intimacies: Influencers and perceived interconnectedness". *Ada: A Journal of Gender, New Media, and Technology*. 8. https://doi.org/10.7264/N3MW2FG2.

Blei, D.M.; Ng, A.Y. & Jordan, M.I. (2003). "Latent Dirichlet allocation". *Journal of Machine Learning Research*. 3: 993-1022. https://doi.org/10.1162/jmlr.2003.3.4-5.993.

Blum-Ross, A. & Livingstone, S. (2017). "'Sharenting': Parent blogging and the boundaries of the digital self". *Popular Communication*. 15(2): 110-125. https://doi.org/10.1080/15405702.2016.1223300.

Cohen, S. (2011). *Folk Devils and Moral Panics: The Creation of the Mods and Rockers*. Routledge Classics ed.

Dal, A. & Nisbet, E.C. (2022). "Walking through firewalls: Circumventing censorship of social media and online content in a networked authoritarian context". *Social Media + Society*. 8(4). https://doi.org/10.1177/20563051221137738.

Divon, T.; Annabell, T. & Goanta, C. (2026). "Children as concealed commodities: Ethnographic nuances and legal implications of kidfluencers' monetisation on TikTok". *New Media & Society*. https://doi.org/10.1177/14614448241304657.

Entman, R.M. (1993). "Framing: Toward clarification of a fractured paradigm". *Journal of Communication*. 43(4): 51-58. https://doi.org/10.1111/j.1460-2466.1993.tb01304.x.

Feller, G. & Burroughs, B. (2021). "Branding kidfluencers: Regulating content and advertising on YouTube". *Television & New Media*. https://doi.org/10.1177/15274764211052882.

Greyson, D.; Chabot, C.; Mniszak, C. & Shoveller, J.A. (2023). "Social media and online safety practices of young parents". *Journal of Information Science.* 49(5): 1344-1357. https://doi.org/10.1177/01655515211053808.

Helsper, E.J. & Eynon, R. (2013). "Digital natives: Where is the evidence?". *British Educational Research Journal.* 36(3): 503-520. https://doi.org/10.1080/01411920902989227.

Kermani, H. (2023). "#MahsaAmini: Iranian Twitter activism in times of computational propaganda". *Social Movement Studies*. https://doi.org/10.1080/14742837.2023.2180354.

Kermani, H. & Tafreshi, A. (2022). "Walking with Bourdieu into Twitter communities: An analysis of networked publics struggling on power in Iranian Twittersphere". *Information, Communication & Society*. https://doi.org/10.1080/1369118X.2021.2021267.

Lee, D.D. & Seung, H.S. (1999). "Learning the parts of objects by non-negative matrix factorization". *Nature*. 401(6755): 788-791. https://doi.org/10.1038/44565.

Livingstone, S. & Helsper, E.J. (2013). "Children, internet and risk in comparative perspective". *Journal of Children and Media.* 7(1): 1-8. https://doi.org/10.1080/17482798.2012.739751.

Livingstone, S.; Mascheroni, G. & Staksrud, E. (2018). "European research on children's internet use: Assessing the past and anticipating the future". *New Media & Society.* 20(3): 1103-1122. https://doi.org/10.1177/1461444816685930.

Livingstone, S. & Stoilova, M. (2021). "The 4Cs: Classifying online risk to children" (CO:RE Short Report Series on Key Topics). Leibniz-Institut für Medienforschung | Hans-Bredow-Institut. https://doi.org/10.21241/ssoar.71817.

Marwick, A.E. (2008). "To catch a predator? The MySpace moral panic". *First Monday.* 13(6). https://doi.org/10.5210/fm.v13i6.2152.

Mendelsohn, J.; Budak, C. & Jurgens, D. (2021). "Modeling framing in immigration discourse on social media". In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (pp. 2219-2263). Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.naacl-main.179.

Prensky, M. (2001). "Digital natives, digital immigrants: Part 1". *On the Horizon.* 9(5): 1-6. https://doi.org/10.1108/10748120110424816.

Shokrollahi, O.; Hashemi, N. & Dehghani, M. (2021). "Discourse analysis of Covid-19 in Persian Twitter social networks using graph mining and natural language processing". *arXiv*. https://arxiv.org/abs/2109.00298.

van Atteveldt, W. & Peng, T.Q. (2018). "When communication meets computation: Opportunities, challenges, and pitfalls in computational communication science". *Communication Methods and Measures*. 12(2-3): 81-92. https://doi.org/10.1080/19312458.2018.1458084.