Posting Time, Relative Network Conditions and Engagement

Evidence from 2.6 Billion TikTok Posts Matched to Hourly National Internet Measurements

KASPR Research

Working paper, October 2026

Abstract

We match 2.58 billion TikTok posts published in 236 countries between 2023 and 2025 to hourly national latency and jitter measurements. About 85–86% of posts in each year are US-labelled, and the United States supplies 78.2% of the identifying information for standardized latency. The pooled estimates therefore primarily describe this US-centred sample, with generalisability to other markets assessed through the country-sensitivity comparisons. Both internet variables are expressed relative to each country's own year. The main model controls for country-specific daily shocks, recurring country-specific hour-of-week patterns, post characteristics and individual snapshot age. Standardized latency is negatively associated with views, likes, comments and shares; jitter is positively associated with engagement conditional on latency. The likes association has stronger statistical support across the reported checks, while the shares association varies more across specifications. Equal-country and non-US estimates are less precise. Separate descriptive models estimate posting-hour and weekday profiles with the same standardized inputs. Levels, logs, clipped measures and bins are retained as robustness checks; the standardized presentation was selected after this exploratory comparison. The findings describe associations in the observed sample and provide a basis for further study of network conditions and engagement.

1 Introduction

Early engagement helps determine how widely a platform distributes a short video (Cheng et al., 2014; Szabo and Huberman, 2010). The audience available when a post appears and the network conditions it faces may therefore relate to its eventual reach and interactions. Creators can choose when to publish, but observational differences between posting slots need not describe what would happen if the same post were moved.

Research on video quality finds that start-up delay and rebuffering affect viewing (Dobrian et al., 2011; Krishnan and Sitaraman, 2013; Nam et al., 2016). Research on social media documents daily and weekly activity patterns, competition for attention and the predictive value of publication timing (Golder and Macy, 2011; Grinberg et al., 2013; Spasojevic et al., 2015; Weng et al., 2012). We combine post metadata with independent national network measurements to examine their relationship on a common hourly panel.

We use country-year standardized latency and jitter as the main network measures. This targets the association with conditions that are unusual for a country in a given year, rather than imposing one common effect for a millisecond of delay across all markets. The distinction matters because absolute probe latency and its variability differ across countries. We retain the other transformations as robustness checks and make the timing of this presentation choice explicit.

The analysis separates a network specification with rich country-specific time controls from descriptive posting-time profiles. Both use the same standardized internet variables and exact individual-age adjustment. The estimates are conditional associations. The release is US-centred:

about 85–86% of posts in each year are US-labelled, and the US supplies 78.2% of the coefficient-specific information for standardized latency. We therefore interpret the pooled results primarily in relation to this sample. Country-weighting and non-US comparisons show how far the findings extend across markets; collection-period, time-control, temporal-exposure and functional-form comparisons provide further context.

2 Related Work

Quality of experience. Across content providers, Dobrian et al. (2011) find that buffering has the largest effect on play time. Using quasi-experimental matching on delivery data, Krishnan and Sitaraman (2013) find that viewers begin abandoning a video after about two seconds of start-up delay. Balachandran et al. (2013) model quality and engagement jointly, while Hoßfeld et al. (2012) and Seufert et al. (2015) relate initial delay and stalls to perceived quality. Nam et al. (2016) find that perceived quality predicts abandonment better than network-level measures. These studies observe individual streams. We observe national network conditions and post-level outcomes, allowing us to examine whether stream-level findings have a counterpart in aggregate engagement.

Latency and online behaviour. Controlled experiments at web search and retail firms find that added hundreds of milliseconds reduce searches, clicks and revenue (Brutlag, 2009; Kohavi et al., 2009), consistent with perceptual thresholds (Nielsen, 1993). Those experiments manipulate a service's own response time; we use natural hour-to-hour variation of national latency.

Timing and attention on social media. Daily and weekly activity cycles appear across cultures (Golder and Macy, 2011; Grinberg et al., 2013). Posts compete for limited attention (Weng et al., 2012), and prior work has studied posting-time recommendations as a prediction problem (Spasojevic et al., 2015). Early engagement also helps predict eventual reach (Cheng et al., 2014; Lerman and Hogg, 2010; Szabo and Huberman, 2010), giving a reason to examine conditions at publication. We extend this work with cross-national estimates that account for network conditions and country-specific seasonal patterns. Following Hofman et al. (2017), we distinguish predictive usefulness from explanation.

Internet measurement and inference. Our network data use active measurement (Bajpai and Schönwälder, 2015). Prior studies document diurnal congestion in access networks (Sundaresan et al., 2011) and show that national aggregates respond to large events (Dainotti et al., 2011). We estimate high-dimensional fixed-effects models (Bergé, 2018; Correia, 2016) and report cluster-robust and multiway-clustered inference (Cameron and Miller, 2015; Cameron et al., 2011). Statistical significance alone offers limited reassurance when signals are weak: estimates can have inflated magnitudes or the wrong sign (Gelman and Carlin, 2014). Undocumented sampling and other limitations of large observational platform datasets also require care (Lazer et al., 2014; Olteanu et al., 2019).

3 Data

3.1 TikTok posts

The datasocial release contains metadata for 5.6 billion TikTok videos in monthly files, collected through TikTok's mobile application programming interface (datasocial, 2026). No sampling design is published. For 2023–2025 we read 17 fields of every post: identifier, posting time (UTC), time of the engagement snapshot, country code, language, type (video or photo post), duration, image count, paid-promotion and shop flags, and the counts of views, likes, comments, shares,

saves and downloads together with the release's engagement-rate field. Within-month identifier uniqueness was checked for the 2025 extraction; global uniqueness across the three-year release was not revalidated here. Deleted posts cannot be observed, and the release carries no account identifier, so creators cannot be followed across posts and spam or automated accounts cannot be identified.

Country. The country field is the two-letter code attached to each post by the platform's interface. The release does not document whether it refers to the account's registered region or the location at upload; it does not refer to the audience. We upper-case two-letter codes, map them to ISO 3166 alpha-3, and drop the 0.01% of rows with other content. 85–86% of posts in each year carry the code of the United States; the ten largest other countries (Indonesia, Pakistan, the Philippines, Mexico, the United Kingdom, Nigeria, Brazil, Vietnam, Germany, Belarus) account for 6–7%. Geographic coverage is broad, but the post-weighted estimates primarily reflect US-labelled posts. This post share is distinct from the US share of identifying information for a particular coefficient, which is 78.2% for standardized latency.

Engagement snapshots. Counts are cumulative at a single snapshot. Posts from 2023, 2024 and January–July 2025 were collected from an archive with snapshots from 31 July 2025 onwards; later posts were collected within weeks. The post-weighted median cell age index is 990 days for 2023 posts, 593 days for 2024 and 112 days for 2025 (Table 1). Historical models control for cell-mean log age; the primary revised model controls for individual log age.

3.2 Network measurements

The network platform sends probes to IPv4 addresses worldwide and records their round-trip times (KASPR, 2026). For each country and UTC hour, it reports the mean round-trip time of responding addresses (latency), its standard deviation (jitter), the number of responding addresses and a flag for regular measurement days. Latency and jitter are measured in milliseconds. Median national latency is about 200 ms, far above a typical broadband connection's latency to a nearby server. The measure reflects national network conditions as seen by external probes; it does not measure the connection experienced by an individual viewer.

Reliable hours. We retain regular measurement hours with at least 30 responding addresses and positive latency and jitter. The analysis sample contains 2,579,434,940 posts in 16,005,925 cells across 236 countries. The sample and reference distributions remain fixed across the functional-form comparisons.

Relative conditions. Both internet measures are standardized within country and UTC calendar year. A value of zero denotes that country's annual mean and a value of one denotes one annual standard deviation above it. The reference distribution gives each distinct reliable country-hour represented in the full analysis sample equal weight, rather than weighting hours by their number of posts. Figure 1 illustrates the resulting exposure series.

3.3 Matching and local time

Each post is matched to the network measurement of its country and the UTC hour that contains its posting time; 99.9% of posts match. Local hour and day of the week are derived from one IANA time zone per country with daylight-saving time resolved for each date. For countries spanning several time zones this is an approximation: the United States is assigned Central Time, so posts from the East Coast are assigned one hour too early and posts from the West Coast two hours too late. This approximation can blur or shift the estimated US profile, depending on

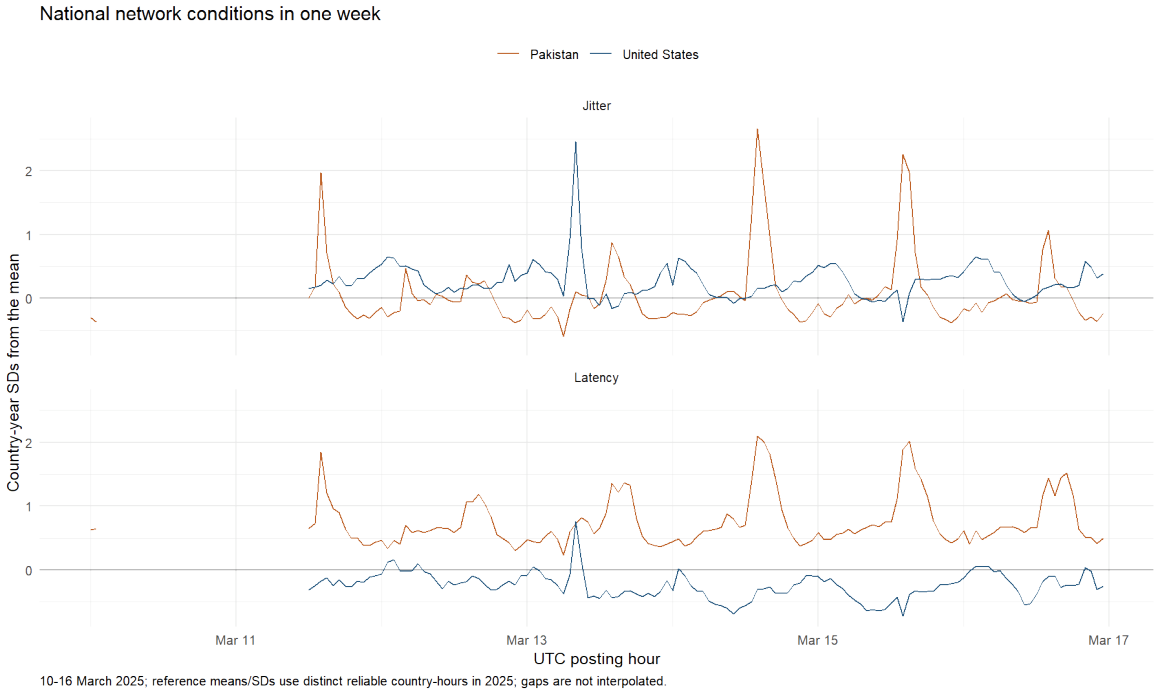


Figure 1: Country-year standardized national latency and jitter, 10–16 March 2025 (UTC): United States and Pakistan. Values are deviations from each country's 2025 mean in units of its 2025 standard deviation.

the geographic mix of posts. A single time zone avoids that particular approximation in other countries but does not establish audience location.

4 Empirical Strategy

4.1 Country-year standardized network measures

For country c, UTC posting hour t and UTC calendar year $y(t)$, define

$$Z^L_{ct} = \frac{L_{ct} - \bar{L}_{cy(t)}}{s^L_{cy(t)}}, \qquad Z^J_{ct} = \frac{J_{ct} - \bar{J}_{cy(t)}}{s^J_{cy(t)}}. \tag{1}$$

The means and sample standard deviations are computed over distinct reliable country-hours represented in the full analysis sample. They are not computed over repeated post cells or weighted by post volume. We retain this reference distribution when changing sample or weights. A country-year with zero or undefined dispersion receives a zero standardized regressor for that measure and supplies no variation in it.

One unit is one country-year standard deviation, not a common millisecond change. This common-effect restriction differs from a pooled slope in absolute units. Standardization also changes countries' relative contribution to estimation; it does not make their contributions equal.

4.2 Network specification

For each of views, likes, comments, shares and saves, we estimate

$$\begin{aligned}\log(1 + Y_{ict}) = {} & \beta_L Z^L_{ct} + \beta_J Z^J_{ct} + \gamma \log(1 + A_i) + \theta \, \mathrm{ad}_i \\ & + \alpha_{c,d(t)} + \eta_{c,24w(c,t)+h(c,t)} + \delta_{b(i)} + \psi_{s(i)} + \varepsilon_{ict}.\end{aligned} \tag{2}$$

Here A_i is individual snapshot age in days, $b(i)$ the duration/type bin and $s(i)$ shop status. Age is calculated in completed hours divided by 24 and truncated below zero. The fixed effects absorb

Table 1: Data coverage, by year of posting.

	2023	2024	2025
posts in the release	524,034,977	693,127,959	1,473,371,731
matched to a country-hour measurement	523,546,277	692,606,643	1,472,579,791
in reliable hours (analysis sample)	500,519,979	661,901,600	1,417,013,361
countries	235	234	234
share United States	85.6%	84.7%	86.2%
share ten largest other countries	6.5%	6.9%	6.4%
share photo posts	5.5%	11.4%	9.4%
share paid (Spark Ads)	0.1%	0.4%	2.0%
cell-mean age index (days):			
weighted p10 / p50 / p90	838 / 990 / 1153	422 / 593 / 779	38 / 112 / 339
mean national latency (ms), post-weighted	200	203	199

Analysis sample: 2,579,434,940 posts in 16,005,925 country × hour × post type × ad × duration cells, 236 countries, 158 weeks. The age index is $\exp(\log(1 + A_i)) - 1$; its percentiles are post-weighted across cells, rather than quantiles of individual ages.

additive country-specific UTC-day shocks, recurring country-specific local hour-of-week patterns, duration/type composition and shop status. All posts receive equal weight in the main model.

The coefficients use remaining within-day exposure variation after these controls. Unusual audience activity and the mix of creators or content can still vary within a day and correlate with network conditions. Stronger time effects also remove genuine slower-frequency network variation. Neither these controls nor a change in functional form establishes exogenous exposure.

4.3 Descriptive posting-time profiles

The country-hour-of-week effects in model (2) absorb recurring posting rhythms separately by country. To describe a common profile, we fit a separate model with country-by-calendar-week effects, global local-hour and weekday indicators, and the same standardized network, post-characteristic and individual-age controls. The reference slots are 16:00 local time and Monday. The hour and weekday components are additive; we do not interpret them as a causal scheduling rule.

4.4 Aggregation, outcome scale and inference

We collapse posts into country × UTC hour × type × ad × shop × duration-bin cells, retaining post counts and means of $\log(1 + Y)$. Individual snapshot age varies within cells. We additionally retain its within-cell sums of squares and its cross-products with each outcome, which recover the individual-age least-squares normal equations and cluster scores exactly. Appendix B gives the construction. DuckDB performs aggregation (Raasveldt and Mühleisen, 2019); weighted demeaning uses `fixest` (Bergé, 2018).

We report 100β in log points. The outcome is mean log(1+count); these numbers are not percentage changes in arithmetic mean counts. The distinction is particularly relevant for sparse interactions and zero counts.

Our main inference clusters two ways, by country and UTC date, allowing within-country dependence across days and common date shocks (Cameron and Miller, 2015; Cameron et al., 2011). We also retain country-only and country-day inference. We apply $G/(G-1)$ to each covariance component, omit a correction based on the number of cell rows, and use a t reference distribution with the minimum cluster count minus one degrees of freedom. Confidence intervals are individual 95% intervals. Holm p-values adjust the five outcomes separately for each network measure.

We chose the standardized representation for the main presentation after comparing seven

functional forms. Those comparisons were exploratory. We retain every form and the additional adjustment across all 70 original form/measure/outcome tests, so the selected model is not presented as independent confirmatory evidence. A large nominal cluster count also offers limited reassurance when a few countries supply most of the information (Gelman and Carlin, 2014).

4.5 Robustness and temporal comparisons

We compare country-week with country-day controls, equal-country with equal-post weights, the sample excluding the United States, and the live-collected August–December 2025 sample. All these comparisons use the original full-sample country-year means and standard deviations.

The joint temporal specification retains contemporaneous latency and jitter and adds latency one day before, one day after and seven days after publication, with the same post controls and fixed effects as the main model. Each exposure is normalized using the posting country-year reference distribution, including at year boundaries. Including current latency is relevant given persistence in network conditions: the future coefficients measure associations conditional on current latency and the other included exposures. The baseline and joint models use identical observations with reliable exposure at every required time. A separate future-only comparator replaces current latency with latency seven days ahead; it is reported as a distinct comparison. Because engagement accumulates after publication, we interpret future-exposure estimates as temporal diagnostics. They are not event-study pre-trends.

5 Results

5.1 Standardized network associations

Table 2 reports the main specification. A coefficient is the adjusted difference in mean $\log(1 + Y)$ associated with a one-standard-deviation change relative to that country's own year. Latency and jitter enter jointly. These are not percentage changes in arithmetic mean counts. Latency

Table 2: Main country-year standardized specification. Estimates are log points per country-year SD, with individual 95% country/date intervals. Holm p adjusts five outcomes separately for each measure.

Outcome	Latency [95% CI]	Holm p	Jitter [95% CI]	Holm p
comments	-0.077 [-0.136, -0.017]	0.023	+0.084 [+0.052, +0.115]	7.00e-07
likes	-0.389 [-0.501, -0.276]	2.94e-10	+0.498 [+0.403, +0.592]	1.57e-20
saves	-0.012 [-0.062, +0.037]	0.625	+0.037 [+0.003, +0.070]	0.032
shares	-0.073 [-0.127, -0.020]	0.023	+0.119 [+0.085, +0.153]	1.34e-10
views	-0.194 [-0.248, -0.140]	7.59e-11	+0.284 [+0.240, +0.328]	6.99e-28

is negatively associated with views, likes, comments and shares, while its saves coefficient is imprecise. Jitter is positively associated with these outcomes conditional on latency. Latency and jitter therefore describe distinct conditional relationships. With the broader Holm adjustment across the original 70 exploratory tests, the negative latency associations for views and likes remain significant at the 5% level; comments and shares do not. This broader family is retained because the standardized form was selected after comparison with the alternatives. The United States supplies 78.2% of the coefficient-specific information for standardized latency, alongside its roughly 85–86% share of posts in each year. These are distinct measures of the sample's US concentration. The pooled estimates are most informative about this US-centred collection; the country comparisons below provide context for generalisability. The inverse-Herfindahl information diagnostic corresponds to 1.6 equally informative countries and describes concentration rather than inferential degrees of freedom.

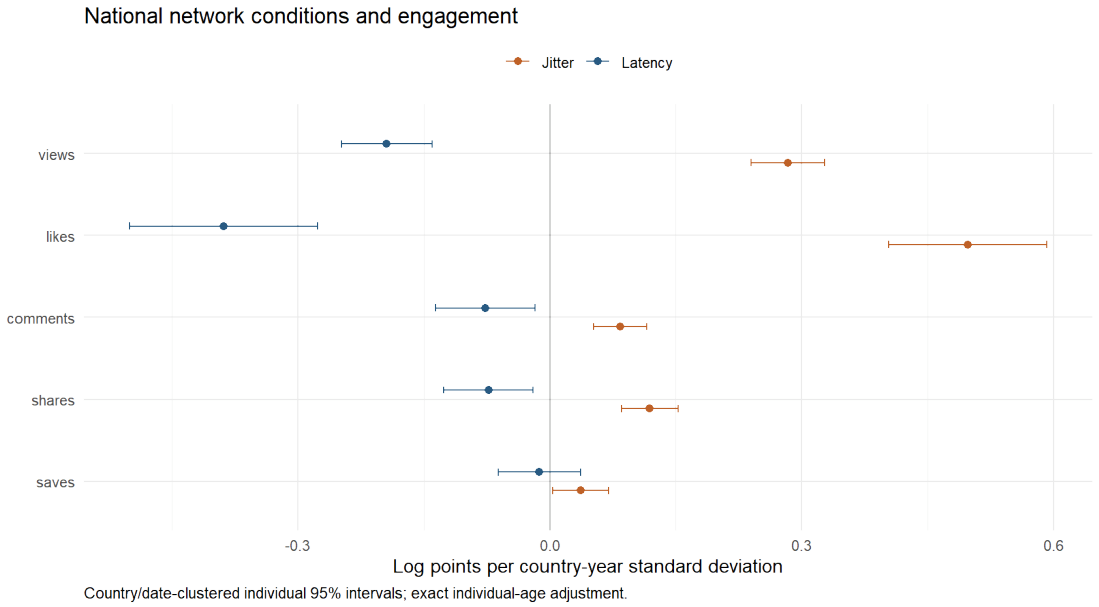


Figure 2: Main standardized latency and jitter associations, with individual 95% country/date intervals.

5.2 Posting-time profiles

The separate country-week timing model includes the same standardized network measures and individual-age adjustment. Figures 3 and 4 describe fitted differences between posts published at different local hours and weekdays. They do not estimate the gain from moving the same post. For this timing model, the two-way covariance requires truncating negative eigenvalues for every outcome; its reported intervals use that adjustment. Country-only covariance is retained in the replication files.

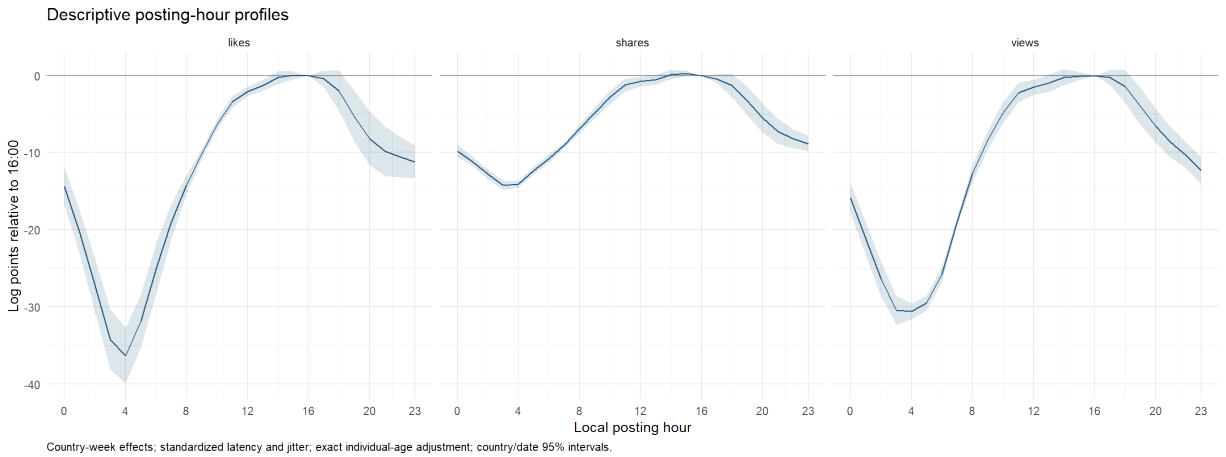


Figure 3: Posting-hour profiles relative to 16:00 local time, in log points; standardized network controls throughout.

5.3 Country weighting, samples and fixed effects

We retain the same country-year reference means and standard deviations in every sample and weighting check. Differences therefore reflect the observations, weighting or controls rather than a redefinition of one standard deviation. Likes has stronger support across the reported checks, while shares varies more with the specification. Both estimates are less precise with equal-country

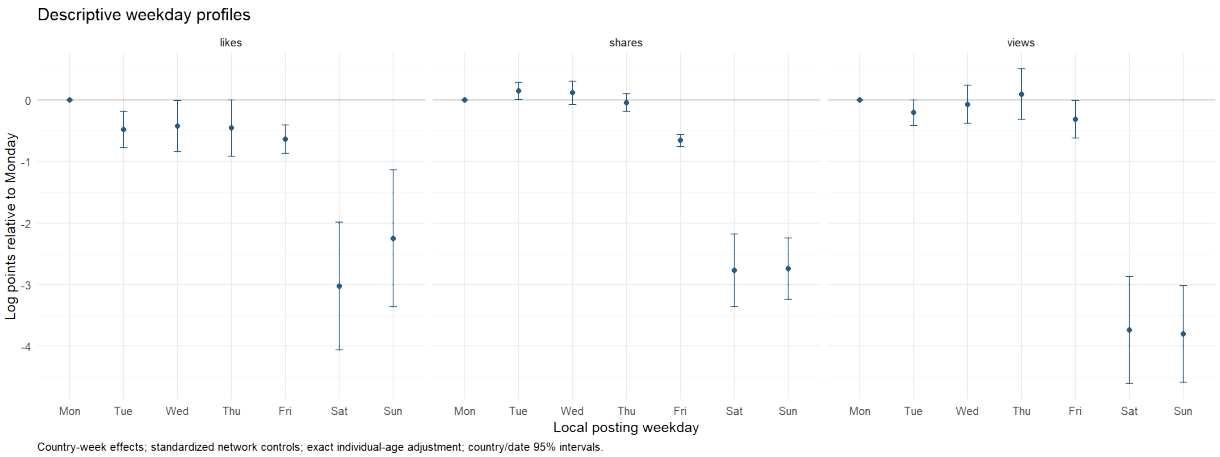


Figure 4: Posting-weekday profiles relative to Monday, in log points. Both profiles use country-week effects and individual 95% country/date intervals.

weights or without the US, which leaves uncertainty about their magnitude and direction across other markets.

Table 3: Standardized sample and fixed-effect checks: log points per country-year SD, with 95% country/date intervals. Other outcomes and country-only inference are retained in the coefficient files.

	Latency		Jitter	
Sample/model	Likes	Shares	Likes	Shares
Equal country weights	-0.108 [-0.600, +0.384]	+0.054 [-0.358, +0.467]	+0.203 [-0.126, +0.532]	+0.004 [-0.235, +0.243]
Country-week effects	-0.261 [-0.396, -0.125]	+0.012 [-0.033, +0.058]	+0.326 [+0.237, +0.414]	+0.065 [+0.042, +0.088]
Live collection only	-0.185 [-0.437, +0.066]	+0.016 [-0.062, +0.094]	-0.047 [-0.249, +0.155]	-0.063 [-0.174, +0.047]
Main, post-weighted	-0.389 [-0.501, -0.276]	-0.073 [-0.127, -0.020]	+0.498 [+0.403, +0.592]	+0.119 [+0.085, +0.153]
Exclude United States	-0.157 [-0.383, +0.069]	+0.006 [-0.167, +0.180]	+0.238 [+0.035, +0.442]	+0.087 [-0.089, +0.263]

5.4 Temporal specificity

The joint temporal model includes contemporaneous latency, latency one day before, one day after and seven days after publication, together with contemporaneous jitter and the main controls and fixed effects. Retaining current latency allows the future-exposure coefficients to be interpreted conditional on posting-hour conditions, which is relevant given temporal persistence. All exposures use the posting country-year reference distribution. The baseline and extended models use identical observations with reliable exposure at each time. With cumulative engagement snapshots, these estimates describe temporal associations around publication. All lead and lag coefficients and

Table 4: Standardized temporal comparison on identical observations. Log points per posting-country-year SD; 95% country/date intervals. The extended model retains current latency and adds all three prior/future exposures.

Outcome	Baseline current	Extended current	Extended +7 days
comments	-0.048 [-0.100, +0.004]	-0.057 [-0.106, -0.009]	+0.001 [-0.026, +0.028]
likes	-0.373 [-0.487, -0.259]	-0.418 [-0.544, -0.292]	+0.094 [+0.043, +0.144]
saves	+0.014 [-0.031, +0.058]	+0.001 [-0.041, +0.043]	-0.004 [-0.032, +0.023]
shares	-0.045 [-0.090, -0.001]	-0.061 [-0.103, -0.019]	-0.003 [-0.033, +0.026]
views	-0.134 [-0.178, -0.090]	-0.171 [-0.214, -0.128]	+0.064 [+0.001, +0.126]

joint tests are retained in the replication files. The separate future-only comparator replaces current latency with latency seven days ahead; it is distinct from the joint model shown here. For likes, future latency remains associated with engagement after conditioning on current latency. The temporal comparisons help describe the exposure window, with causal timing left for further

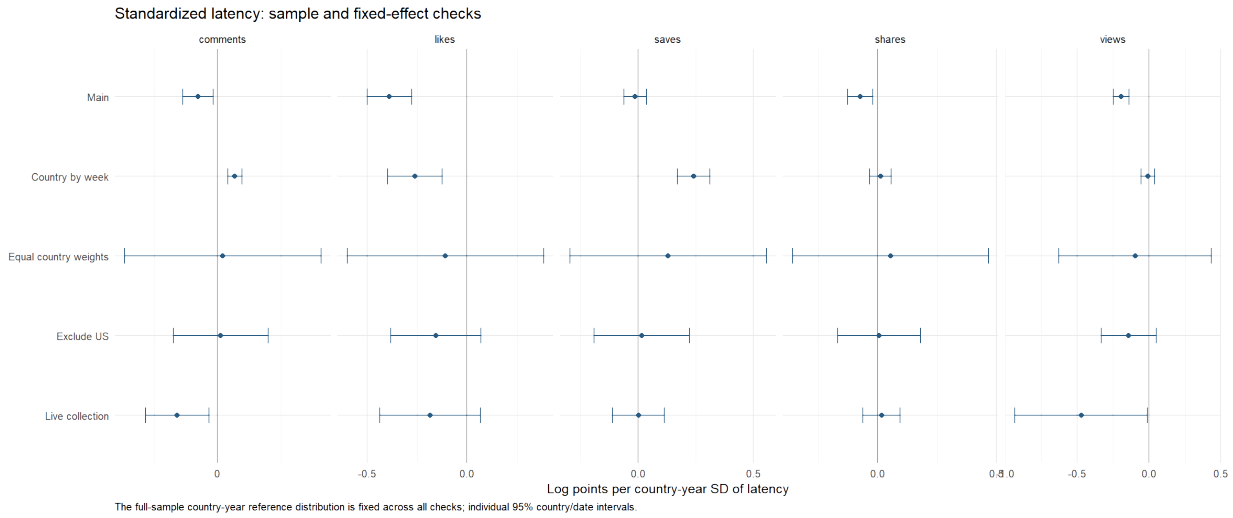


Figure 5: All five outcomes under the standardized sample and fixed-effect checks.

study. Some two-way covariance estimates require truncating negative eigenvalues; those cases are flagged in the files. Country-only inference is retained as a sensitivity check.

6 Discussion

The standardized specification examines how engagement differs when national network conditions are unusual for a country's own year. Latency and jitter enter jointly, so their negative and positive coefficients describe distinct conditional relationships. Jitter can covary with load, routes and the composition of responding addresses, as well as with the mix of posts and viewers. The two measures therefore warrant separate interpretation.

The findings are best understood as evidence from a large, US-centred collection. US-labelled posts account for about 85–86% of observations in each year, and the US supplies 78.2% of the coefficient-specific information for standardized latency. Coverage of 236 countries adds useful geographic breadth, while the pooled estimates remain most informative about the sample's largest market. Generalisability depends on how other markets compare in audience behaviour, platform use and network conditions. The equal-country and non-US estimates for likes and shares are less precise, leaving more uncertainty about the magnitude and direction of those relationships outside the main sample composition.

The negative latency association for likes has stronger support across the reported specifications and multiplicity adjustments. The shares association is more sensitive to the choice of controls, weights and functional form. Quartile profiles also suggest that the relationship need not be uniformly decreasing across the exposure distribution. These comparisons help define the scope of the pooled findings.

The temporal diagnostics include current latency alongside the prior and future exposures. For likes, future latency remains associated with engagement after conditioning on posting-hour latency in the joint model. Given cumulative engagement snapshots, these estimates are consistent with a relationship involving network conditions over the period around publication; they do not by themselves isolate its timing or mechanism.

The posting-time profiles compare different posts published in different slots. They combine audience availability with content and creator selection, competing supply and other conditions that remain unobserved. We have not evaluated a publication-time intervention or a forecast on held-out posts. The full-year normalization is retrospective; a deployable system would require a reference distribution estimated from information available before publication and a separate evaluation of its recommendations.

7 Limitations

The release has an undocumented sampling design and no account identifier, which limits adjustment for creator characteristics. The country code describes the post's recorded region, while audience location is unavailable. National probe round-trip times provide an aggregate measure of network conditions, and one time zone per country approximates posting time in multi-zone countries. These features should be considered when extending the results to particular audiences or settings.

Engagement is observed at one cumulative snapshot. The exact individual-age control accounts for age variation within cells, while differences in collection timing and the archive/live transition remain relevant to interpretation. Unobserved within-day differences in posts and audiences may also contribute to the associations. We therefore interpret the estimates descriptively. Some multiway covariance estimates require eigenvalue truncation, which is flagged in the outputs; country-only results remain available.

The main transformation was selected after comparing alternative forms. We retain the full comparison and the broader multiple-testing adjustment so readers can assess this exploratory choice. The country-sensitivity results provide context for generalisability, and independent samples would help establish how well the patterns extend to other settings.

8 Conclusion

Across 2.6 billion TikTok posts, country-year standardized national latency is associated with modestly lower views and interactions in the post-weighted model, while jitter is positively associated with engagement conditional on latency. These findings primarily describe a US-centred sample: about 85–86% of posts in each year are US-labelled, and the US contributes 78.2% of the identifying information for latency. The likes association has stronger statistical support across the reported checks; the shares association warrants a more qualified interpretation. The country-sensitivity comparisons indicate uncertainty about the size and direction of these relationships in other markets. The descriptive posting-time profiles complement these findings, with audience-relevant exposure and engagement measured over a fixed early horizon offering useful directions for future work.

Data Availability and Reproducibility

The TikTok release is available under CC BY-NC 4.0 (datasocial, 2026); national network measurements are proprietary and available from the authors on request for replication. The current entry point is `analysis/R/28_standardized_main.R`. Reports and manuscript outputs are generated by `analysis/R/29_standardized_report.R`. Exact individual-age replication requires the within-cell age moments reconstructed from all 36 raw monthly files, as well as the original cell-level covariates and the fixed country-year reference distribution. Other transformations and their original comparison remain available with the analysis.

Ethics Statement

The analysis uses released post metadata and national network aggregates, without text, images, audio or account identifiers. We report aggregates and do not attempt creator or viewer identification. Use follows the release's non-commercial licence.

References

Vaibhav Bajpai and Jürgen Schönwälder. A survey on internet performance measurement platforms and related standardization efforts. *IEEE Communications Surveys & Tutorials*, 17(3):1313–1341, 2015.

Athula Balachandran, Vyas Sekar, Aditya Akella, Srinivasan Seshan, Ion Stoica, and Hui Zhang. Developing a predictive model of quality of experience for internet video. In *Proceedings of the ACM SIGCOMM 2013 Conference*, pages 339–350. ACM, 2013.

Laurent Bergé. Efficient estimation of maximum likelihood models with multiple fixed-effects: The R package FENmlm. *CREA Discussion Papers*, (13), 2018.

Jake Brutlag. Speed matters for Google web search. Technical report, Google, 2009.

A. Colin Cameron and Douglas L. Miller. A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372, 2015.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249, 2011.

Justin Cheng, Lada Adamic, P. Alex Dow, Jon M. Kleinberg, and Jure Leskovec. Can cascades be predicted? In *Proceedings of the 23rd International Conference on World Wide Web*, pages 925–936, 2014.

Sergio Correia. Linear models with high-dimensional fixed effects: An efficient and feasible estimator. Technical report, Duke University, 2016. Working paper.

Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapè. Analysis of country-wide internet outages caused by censorship. In *Proceedings of the 2011 ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 1–18. ACM, 2011.

datasocial. TikTok videos (5.6 billion): One parquet file per month. Hugging Face dataset `datasocial/tiktok-5.6B-videos`, CC BY-NC 4.0, 2026. Accessed 2026-10-13.

Florin Dobrian, Vyas Sekar, Asad Awan, Ion Stoica, Dilip Joseph, Aditya Ganjam, Jibin Zhan, and Hui Zhang. Understanding the impact of video quality on user engagement. In *Proceedings of the ACM SIGCOMM 2011 Conference*, pages 362–373. ACM, 2011.

Andrew Gelman and John Carlin. Beyond power calculations: Assessing type S (sign) and type M (magnitude) errors. *Perspectives on Psychological Science*, 9(6):641–651, 2014.

Scott A. Golder and Michael W. Macy. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881, 2011.

Nir Grinberg, Mor Naaman, Blake Shaw, and Gilad Lotan. Extracting diurnal patterns of real world activity from social media. In *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media*, 2013.

Jake M. Hofman, Amit Sharma, and Duncan J. Watts. Prediction and explanation in social systems. *Science*, 355(6324):486–488, 2017.

Tobias Hoßfeld, Sebastian Egger, Raimund Schatz, Markus Fiedler, Kathrin Masuch, and Charlott Lorentzen. Initial delay vs. interruptions: Between the devil and the deep blue sea. In *Fourth International Workshop on Quality of Multimedia Experience (QoMEX)*, pages 1–6. IEEE, 2012.

KASPR. Hourly national internet connectivity measurements: Methodology. Technical documentation, KASPR, 2026. Data vintage 16 July 2026; available from the authors on request.

Ron Kohavi, Roger Longbotham, Dan Sommerfield, and Randal M. Henne. Controlled experiments on the web: Survey and practical guide. *Data Mining and Knowledge Discovery*, 18(1):140–181, 2009.

S. Shunmuga Krishnan and Ramesh K. Sitaraman. Video stream quality impacts viewer behavior: Inferring causality using quasi-experimental designs. *IEEE/ACM Transactions on Networking*, 21(6):2001–2014, 2013.

David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. The parable of Google Flu: Traps in big data analysis. *Science*, 343(6176):1203–1205, 2014.

Kristina Lerman and Tad Hogg. Using a model of social dynamics to predict popularity of news. In *Proceedings of the 19th International Conference on World Wide Web*, pages 621–630, 2010.

Hyunwoo Nam, Kyung-Hwa Kim, and Henning Schulzrinne. QoE matters more than QoS: Why people stop watching cat videos. In *Proceedings of IEEE INFOCOM 2016*, pages 1–9. IEEE, 2016.

Jakob Nielsen. *Usability Engineering*. Academic Press, 1993.

Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kıcıman. Social data: Biases, methodological pitfalls, and ethical boundaries. In *Frontiers in Big Data*, volume 2, page 13, 2019.

Mark Raasveldt and Hannes Mühleisen. DuckDB: An embeddable analytical database. In *Proceedings of the 2019 International Conference on Management of Data (SIGMOD)*, pages 1981–1984. ACM, 2019.

Michael Seufert, Sebastian Egger, Martin Slanina, Thomas Zinner, Tobias Hoßfeld, and Phuoc Tran-Gia. A survey on quality of experience of HTTP adaptive streaming. *IEEE Communications Surveys & Tutorials*, 17(1):469–492, 2015.

Nemanja Spasojevic, Zhisheng Li, Adithya Rao, and Prantik Bhattacharyya. When-to-post on social networks. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2127–2136, 2015.

Srikanth Sundaresan, Walter de Donato, Nick Feamster, Renata Teixeira, Sam Crawford, and Antonio Pescapè. Broadband internet performance: A view from the gateway. In *Proceedings of the ACM SIGCOMM 2011 Conference*, pages 134–145. ACM, 2011.

Gabor Szabo and Bernardo A. Huberman. Predicting the popularity of online content. *Communications of the ACM*, 53(8):80–88, 2010.

Lilian Weng, Alessandro Flammini, Alessandro Vespignani, and Filippo Menczer. Competition among memes in a world with limited attention. *Scientific Reports*, 2:335, 2012.

A Alternative transformations

We compare the selected standardized form with six alternatives on the identical full sample and with the same exact age adjustment, fixed effects, weights and clustering. The functional-form comparison preceded selection of the main presentation. The full coefficient files retain all five outcomes, both measures and the additional Holm adjustment across the 70 original

Table 5: Latency functional-form robustness. Estimates are log points; effect units vary by row. Holm p adjusts the five outcomes within each form and measure.

Form	Likes [95% CI]	Holm p	Shares [95% CI]	Holm p
Country-year SD (main)	-0.389 [-0.501, -0.276]	2.94e-10	-0.073 [-0.127, -0.020]	0.023
Levels	-0.776 [-1.519, -0.034]	0.121	-0.168 [-0.318, -0.019]	0.108
Log both measures	-2.217 [-4.053, -0.380]	0.073	-0.524 [-0.978, -0.070]	0.073
Log latency	-1.582 [-2.864, -0.301]	0.063	-0.256 [-0.548, +0.035]	0.253
Quartile indicators	-0.080 [-0.246, +0.086]	5.32e-11	+0.157 [+0.082, +0.231]	2.90e-20
Worst decile	-0.697 [-0.973, -0.421]	4.91e-06	-0.021 [-0.059, +0.017]	0.335
Clip p1/p99	-0.575 [-1.219, +0.069]	0.400	-0.137 [-0.323, +0.050]	0.438

form/measure/outcome tests. Levels and clipped levels use a 10 ms contrast; logs use a 10% increase; the standardized model uses one country-year SD; quartiles compare Q4 with Q1; and the worst-decile model compares hours above p90 with the remainder. The quartile p-values test all three bin indicators jointly, whereas their displayed interval concerns Q4 minus Q1. Country-year standardization defines effects relative to each country's own variation and changes countries' relative contribution to estimation. The alternative forms therefore provide complementary comparisons on different exposure scales. Clipping yields less precise negative level slopes. Quartile profiles are non-monotonic for likes, and the highest-latency quartile has more shares than the lowest. These patterns qualify the interpretation of a common linear slope, particularly for shares.

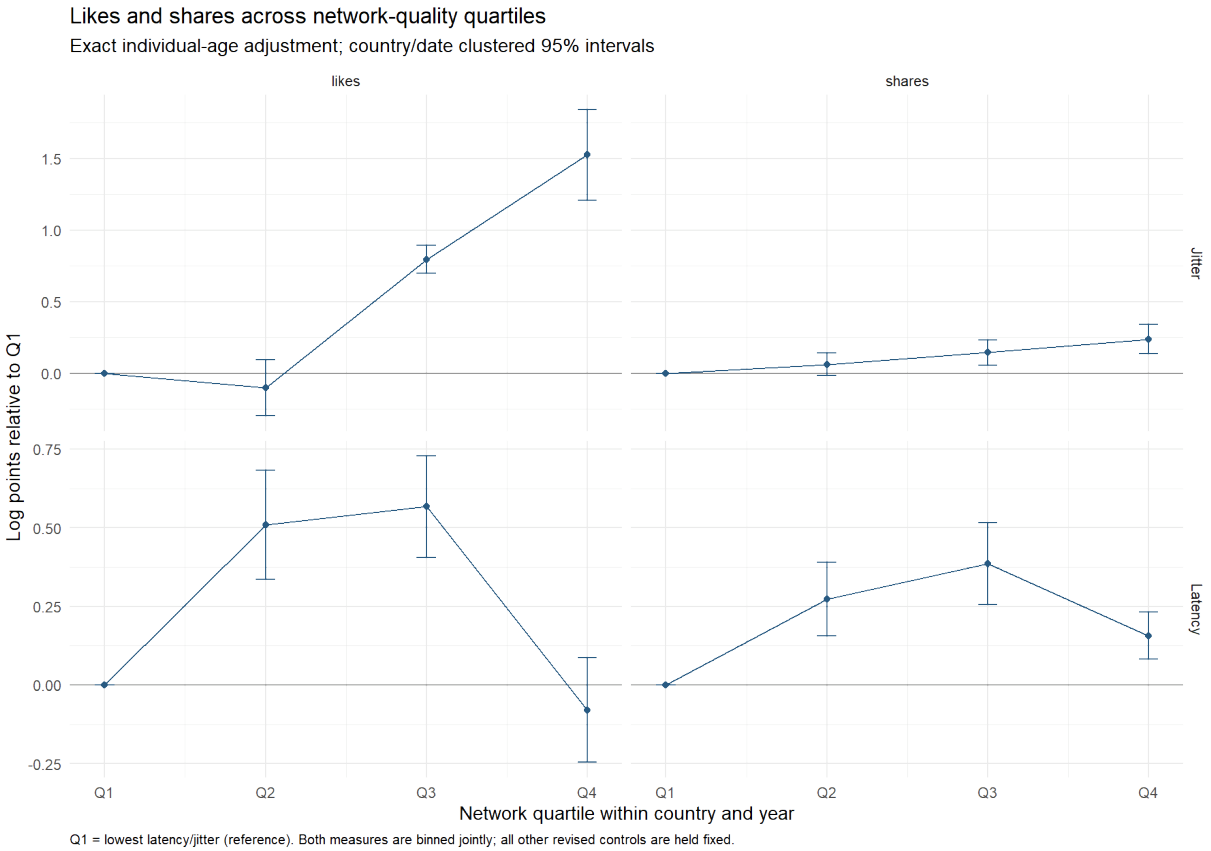


Figure 6: Alternative quartile specification for latency and jitter, with exact individual-age adjustment and country/date intervals. Q1 is the lowest-measure reference quartile.

B Individual-age aggregation

Let g index cells with count n_g, outcome mean $\bar{y}_g$ and cell-constant regressors z_g. For a model containing only those regressors,

$$\sum_{i \in g} (y_i - z_g'\theta)^2 = \sum_{i \in g} (y_i - \bar{y}_g)^2 + n_g(\bar{y}_g - z_g'\theta)^2.$$

This identity preserves point estimates under count weighting. Individual log snapshot age $a_i = \log(1 + A_i)$ is not cell-constant. We retain

$$S_{aa,g} = \sum_{i \in g} (a_i - \bar{a}_g)^2, \qquad S_{ay,g} = \sum_{i \in g} (a_i - \bar{a}_g)(y_i - \bar{y}_g).$$

After weighted demeaning by the fixed effects, add $\sum_g S_{aa,g}$ to the age diagonal of the regressor cross-product matrix and $\sum_g S_{ay,g}$ to the age row of the outcome cross-products. The age score in cluster C additionally receives $\sum_{g \in C}(S_{ay,g} - \hat{\gamma} S_{aa,g})$. For country-balanced weights, each cell's within-post moments receive the corresponding per-post country weight. These moments recover individual-age coefficients and cluster scores for the stated model without materializing billions of rows. Covariance equality additionally requires the same declared finite-sample convention. Synthetic and actual-post comparisons validate the implementation.