

Attention Surges and Recovery in Wikipedia Visits to Korean Music Artists

Frankie Gao

perspectiveopen.com

Correspondence: zcfbfga@ucl.ac.uk

Submitted: 6 July 2026. Published: 6 July 2026.

Abstract

Daily attention peaks are easy to notice, but a falling peak does not establish that attention has returned to its previous level. This study develops a transparent episode description that separates an observable surge alert, confirmed baseline return, residual attention burden and renewed attention. The data comprise 10,404 daily user-classified visits to six English Wikipedia artist pages from 1 January 2022 to 30 September 2026. A past-only detector combines a robust log-count deviation of at least three with a count at least twice the preceding 28-day median. Accepted onsets are separated by 29 days. Recovery requires three consecutive post-onset days at or below 1.25 times the frozen pre-alert median. Follow-up is limited to 28 days, with explicit right censoring. The detector identifies 92 episodes, of which 91 have complete follow-up. The pooled Kaplan–Meier median confirmed return is day 8. Among complete episodes, 42 return by day 7, 64 by day 14 and 78 by day 28. Updating the reference median every day raises the day-28 count to 90, showing that recovery classifications depend materially on reference adaptation. The median positive excess burden is 9.80 baseline equivalent days, and the median share accruing on days 8–28 is 37.5%. Ninety episodes confirm a half-excess decline after the first-week peak, but 12 of those still lack full baseline return by day 28. Renewed attention occurs in 14 of 73 episodes with a full seven-day observation opportunity after return. Rule perturbations and calendar-block resampling quantify descriptive sensitivity. The results support monitoring that records intensity, return status and residual burden separately. They describe page-visit dynamics within a selected historical panel and do not identify unique fans, event causes or promotional effects.

Keywords: collective attention; Korean music; Wikipedia pageviews; surge detection; recovery duration; residual attention; calendar block bootstrap

1 Introduction

Music promotion and audience research frequently focus on the date and height of a visible attention peak. A large peak may justify immediate scrutiny, yet it gives incomplete information about the subsequent monitoring workload. One artist can receive an intense, short-lived increase and quickly return to typical traffic. Another can remain moderately elevated for several weeks. A third can appear to recover and then experience a second increase. These patterns require different descriptions even when the initial alerts look similar. An accessible research problem is therefore to measure when attention increases, how long the elevated period remains observable and how its later trajectory differs from its initial intensity.

Previous work on collective attention has shown that response trajectories can contain useful information beyond their maximum. Crane and Sornette examined heterogeneous activity patterns in a large collection of YouTube videos and linked fitted response classes to a model of social amplification [3]. Lorenz-Spreen and colleagues studied changes in collective-attention dynamics across several domains [4]. These studies

motivate examination of temporal shape, but their mechanistic or population-level conclusions cannot simply be transferred to a small set of Wikipedia pages. The present study uses transparent statistical summaries without assigning endogenous or exogenous causes, estimating a universal decay law or testing whether all music audiences are becoming less attentive.

Wikipedia provides a useful public observation channel because its numerical pageview API offers a consistent daily calendar and openly reusable counts [1,2]. A visit to an artist page may reflect information seeking, a news response, existing interest or repeated reading by the same person. Its value for this study is the observed timing and scale of page traffic. The count is not a direct measure of fandom, sentiment, music consumption or the number of people reached. The selected English-language pages also do not establish the location or nationality of readers. These distinctions are maintained throughout the analysis.

The central measurement difficulty concerns the baseline. An adaptive baseline is valuable for detecting changes against recent conditions. During follow-up, however, an adaptive reference can rise because it includes the very observations whose persistence is being studied. A plateau that remains elevated above the pre-alert level can become ordinary relative to a new median. Both comparisons can be useful, but they answer different questions. This study therefore freezes the baseline for the primary recovery measure and retains an adaptive reference as a paired methodological comparison on the same episodes.

Four questions guide the analysis. First, how often does a fixed past-only rule identify separated attention surges in this panel? Second, what fraction of accepted episodes confirms return within one, two or four weeks? Third, how much positive excess remains after the initial eight days, and how often does a confirmed return precede renewed attention? Fourth, how much do the answers change when the reference, confirmation rule or detector thresholds change? These are descriptive questions about a defined dataset. They are not a claim that an optimal alert threshold has been learned from a labelled population of real events.

The practical contribution is an episode record that keeps several clocks distinct. The alert clock can run using information available at the end of the current day. The recovery clock records an observable completion of a short confirmation sequence. The burden measure sums the later positive deviations from the same historical reference. An additional seven-day check records a renewed increase after recovery. Together, these quantities offer a more informative monitoring record than a peak date alone, while remaining reproducible with basic statistics and a modest local computing environment.

2 Data and Methods

2.1 Data identity and scope

The primary panel contains daily counts for BTS, BLACKPINK, TWICE, Stray Kids, SEVENTEEN and aespa. Each artist has 1,734 dates from 1 January 2022 through 30 September 2026, giving 10,404 artist-date observations. The queried pages are BTS, Blackpink, Twice, Stray_Kids, Seventeen_(South_Korean_band) and Aespa on en.wikipedia.org. Counts use the all-access aggregate, the user agent classification, daily granularity and UTC calendar dates. The user classification is an API traffic category rather than verification that every visit comes from a distinct human. Page-title changes, redirects and other pages about the same artist are not reconstructed.

The source is a preserved Wikimedia API snapshot assembled for an earlier local forecasting analysis. The initial 2024–2026 panel contributed 6,024 observations; the 2022–2023 extension contributed 4,380.

Retrieval occurred on 2 October 2026. This study reuses the original CSV and 12 raw API response files without replacing historical counts or fetching a later revision. Raw responses are decoded independently and reconciled to every CSV count and date. There are no missing artist-date cells, duplicate artist-date keys, negative counts or non-integer counts. No count is imputed. Input hashes and original retrieval provenance accompany the research package.

The numerical API data are distributed under CC0 1.0 according to Wikimedia’s access policy [2]. That statement applies to the numerical data used here. The study does not redistribute Wikipedia article text, portraits or artwork under an assumed equivalent licence. No participant-level records, individual comments, browsing identities or private platform data enter the analysis. The dataset has previously been accessed, so this is a retrospective empirical study rather than preregistration or an untouched prospective test. The Paper 3 rule record was saved before computing the new episode outcomes.

Table 1. Observed pages and panel scale

Artist	Wikipedia page	Days	Median visits	Maximum visits
BTS	BTS	1734	9,433.5	140,144
BLACKPINK	Blackpink	1734	6,159.0	50,338
TWICE	Twice	1734	2,698.0	17,371
Stray Kids	Stray_Kids	1734	3,009.0	20,174
SEVENTEEN	Seventeen_(South_Korean_band)	1734	2,083.0	16,238
aespa	Aespa	1734	1,599.5	15,261

Daily user-classified all-access visits to English Wikipedia. Medians and maxima are over all 1,734 dates per page. Visits can repeat and are not unique audience counts.




Figure 1. Daily visits and accepted surge onsets. Gold points mark the 92 accepted onsets. All artist panels share a logarithmic vertical scale. The first 28 days are baseline warm-up. Point placement identifies an operational alert, without attributing a cause.

2.2 Past only alert rule

For artist i on day t , the count reference B is the median of the preceding 28 counts, excluding the current day. Twenty-eight days cover four complete weekday cycles. The median gives limited influence to isolated large days, although a long elevated period can still shift it. Separately, the detector computes the median and median absolute deviation of $\log(1 + \text{count})$ over the same previous dates. A log scale expresses a relative departure more naturally than a common raw-count difference when the pages have different traffic levels. The median absolute deviation is a robust scale summary; the factor 1.4826 gives the familiar normal-consistent scaling [5].

$$B_{i,t} = \text{median}\{Y_{i,t-28}, \dots, Y_{i,t-1}\} \quad (1)$$


$$Z_{i,t} = \frac{\log(1 + Y_{i,t}) - m_{i,t}}{\max(1.4826 \text{ MAD}_{i,t}, 0.10)} \quad (2)$$

A date is a candidate alert only when its robust score is at least 3 and its count is at least twice $\max(B, 1)$. The 0.10 scale floor prevents an almost constant reference window from generating a very large score from a trivial change. The ratio condition also protects against interpreting a small absolute deviation as a surge. The thresholds are explicit operational choices. A score of three is not a Gaussian significance test, a false-

alert probability or evidence that an external music event occurred. There is no independently labelled event set from which sensitivity, specificity or detection precision could be estimated.

Candidates are processed chronologically within each artist. The first candidate is accepted. Subsequent candidates are accepted only if at least 29 calendar days have elapsed since the preceding accepted onset. This spacing keeps primary day-0-to-day-28 follow-up windows from overlapping within an artist. It does not prevent simultaneous episodes across artists, and it does not imply independence between episodes of the same artist. Suppressed candidates remain visible in the daily detection table. A renewed increase inside the observation window is described as part of that episode rather than counted as a second primary onset.

This design deliberately identifies the first qualifying alert rather than retrospectively moving the onset to the tallest peak. The onset decision can be made after today's UTC count is observed, using the past and current day only. A peak selected from the following week uses later information and is therefore used solely for retrospective trajectory description. The study does not claim that the future peak, return time or event burden could have been known when the alert was first generated.



Alert detection uses past and current observations. Recovery and peaks use later observations.

Figure 2. Information boundaries and episode measurements. All detector inputs come from the preceding 28 days and the current completed day. Follow-up measurements require later dates. The frozen reference is the median immediately before onset, while the adaptive comparison updates its median daily.

2.3 Confirmed return and censoring

The onset reference B is fixed for the primary follow-up. A confirmed return is the first completion day of three consecutive post-onset counts at or below $1.25B$. The sequence may begin on day 1, making day 3 the earliest possible primary confirmation. Duration is measured to the completion day rather than backdated to the first of the three low days. This creates a reproducible information boundary: a monitor cannot confirm a three-day sequence after seeing only its first day. The 25% tolerance permits ordinary variation around the old median without requiring every observation to fall below it.

$$T_{i,e} = \min\{d \geq 3 : Y_{i,t+d-2}, Y_{i,t+d-1}, Y_{i,t+d} \leq 1.25 B_{i,t}\} \quad (3)$$

Every episode is followed through day 28 or the last panel date, whichever comes first. A lack of confirmation is right censored at that observed horizon. The analysis does not label such episodes as lasting exactly 28 days, assign them zero duration or omit them from the return curve. The Kaplan–Meier product-limit estimate retains these incomplete observations [6]. Here the event is confirmed return, so survival denotes the fraction not yet confirmed returned. Risk sets include episodes still unconfirmed and observed immediately before each elapsed day. Confirmations occur before censor removals at a tied day.

$$\hat{S}(d) = \prod_{u \leq d} \left(1 - \frac{D_u}{N_u}\right) \quad (4)$$

The pooled curve includes all 92 primary episodes. Fixed-horizon proportions, total burden and trajectory comparisons use only episodes with the complete 29 observations spanning day 0 through day 28. This restriction ensures an equal observation opportunity for those summaries. The single late-onset episode remains in the censored return analysis but receives no invented tail burden or peak classification. Artist curves are descriptive comparisons of small selected cohorts. No log-rank test, proportional-hazards model or population-level artist ranking is inferred from them.

2.4 Peak decline and residual burden

An early peak is the largest observed count on days 0–6, with the earliest date retained in a tie. Its excess is $P - B$. A confirmed half-excess decline is the first completion of three consecutive counts at or below $B + 0.5(P - B)$, strictly after that early peak and within the observed day-28 horizon. Both the peak choice and subsequent decline are retrospective measurements. The term half-excess decline is used because the trajectories need not follow an exponential decay, can contain later peaks and can cross the same level repeatedly. It is not an estimated physical or behavioural half-life.

Positive excess E on each day is $\max(Y - B, 0)$. Positive excess burden sums E over days 0–28. Dividing by B produces baseline equivalent days. For example, ten baseline equivalent days means that the accumulated positive deviations equal ten days of the pre-alert daily reference count. It does not mean that attention remained continuously elevated for ten days. Counts below the baseline are deliberately excluded from this positive-excess measure rather than offset against earlier increases. The result describes deviations from a historical reference, not additional people or visits caused by a campaign.

$$A_{i,e} = \sum_{d=0}^{28} E_{i,e,d} / B_{i,t} \quad (5)$$

The tail share is the fraction of that positive excess occurring on days 8–28. Its complement corresponds to the initial eight dates, days 0–7. These periods have different lengths, so a large tail share is not automatically evidence of an equally high daily intensity. The analysis also reports the median ratio between the tail’s daily positive excess and the initial period’s daily positive excess. For aligned profiles, each observed signed deviation is divided by the early-peak excess. Negative values and later values above one are retained. Clipping these profiles at zero or one would hide below-baseline counts or later peaks.

2.5 Renewed attention and reference adaptation

Renewed attention is defined as at least one count above $1.5B$ during the seven calendar days after a confirmed return. Only complete episodes whose return is confirmed by day 21 have all seven subsequent dates inside the primary observation window. These episodes form the renewal denominator. A return on

day 25 cannot be treated as a negative renewal because only three of the required seven later days have been observed. The 1.5 threshold is a descriptive renewal flag and does not replace the original alert rule, which additionally requires the robust score and chronological spacing.

The adaptive-reference comparison retains the same 92 accepted onsets and the same three-day confirmation rule, but compares each follow-up count with 1.25 times the count median preceding that particular date. It therefore answers whether traffic has become ordinary relative to the recently updated reference. The primary frozen-reference result instead asks whether traffic has approached the pre-alert level. Paired differences are calculated on the 91 complete episodes to isolate this definition change from changes in the detector or follow-up opportunity. The comparison does not identify one reference as the true causal counterfactual.

2.6 Sensitivity uncertainty and reproducibility

One-factor sensitivity analyses change the reference window to 56 days, the onset ratio to 1.5 or 2.5, the robust-score threshold to 4, the log-scale floor to 0.05 or 0.20, the return tolerance to 1.10 or 1.50, the confirmation length to one day and the onset spacing to 14 or 43 days. The main detector is also summarized for onsets during 2024–2026, while retaining their existing historical reference and spacing state. A 14-day spacing creates overlapping follow-up windows and is explicitly marked as such. Alternatives are reported together rather than selecting the specification with the most favourable conclusion.

To describe sampling sensitivity while respecting calendar dependence, 2,000 circular calendar-block resamples are generated with seed 20261003. The primary block length is 60 days, with 30- and 90-day checks. The resampling calendar comprises all 1,678 dates that can have 28 preceding and 28 following observations. Each date carries the already observed outcome sums and episode count for all six artists jointly, including zero-event dates. Blocks preserve nearby onset clustering and simultaneous artists. The real outcome of each episode stays attached to its original onset record; the detector is not rerun on an artificial concatenation of count series.

For each resample, outcome sums are divided by its sampled episode count. Reported intervals are the 2.5th and 97.5th percentiles for complete-episode return fractions, mean burden, mean tail share and the paired adaptive-minus-frozen day-14 difference. They are conditional descriptive resampling intervals, motivated by block methods for dependent observations [7]. Approximate stationarity over this selected historical period is a limitation, and year-specific changes can weaken the approximation. The intervals do not establish prospective alert performance or uncertainty about unobserved event causes. Equal-artist summaries are additionally calculated so that the pooled results can be checked against disproportionate episode contributions.

Analysis runs locally in Python 3.12.14 using NumPy, pandas, SciPy and Matplotlib. The research package contains a fixed rule record, source provenance, raw API responses, daily candidate and suppression flags, the full episode catalogue, long-format trajectories, numerical tables, all figure sources and an executed notebook. Independent checks reconstruct all 10,236 eligible daily decisions directly from the counts, verify every episode outcome and recompute the 203 pooled and artist survival risk-set rows. Confirmation boundaries are checked with isolated verification fixtures that never enter the empirical dataset. The spreadsheet provides an editable single-episode exploration of the return tolerance while keeping the study's primary results fixed.

3 Results

3.1 Detected episodes and calendar coverage

The 10,236 dates with a full past reference produce 217 candidate alerts. Chronological spacing suppresses 125 candidates and accepts 92 onsets. Each artist has 1,706 eligible detection dates. Accepted episode counts range from 12 for Stray Kids to 18 for BTS and aespa. BLACKPINK has 15 accepted episodes, TWICE 15 and SEVENTEEN 14. A suppressed candidate is a qualifying day inside the spacing period, not an error or evidence that no further increase occurred. The accepted count consequently represents separated monitoring episodes under this rule rather than a census of all music or news events.

Ninety-one episodes have complete follow-up. The BLACKPINK onset of 28 September 2026 has only two later dates before the panel ends on 30 September and is right censored at day 2. It is excluded from fixed 28-day burden and renewal summaries. Of the full panel's 14 censored recovery observations, 13 are unresolved at day 28 and one ends with the panel. This separation prevents the short final observation from being misreported as a long-lasting surge. Figure 3 shows the year distribution of accepted onsets; 2022 loses its first 28 dates to warm-up and 2026 ends in September, so raw annual counts have unequal exposure.




Figure 3. Accepted onsets by artist and year. Cells show counts under the fixed primary rule. The eligible calendar is 337 days in 2022, 365 in 2023, 366 in 2024, 365 in 2025 and 273 in 2026 for each artist. Exposure-adjusted rates are available in the companion table.

Table 2. Detection and follow up accounting

Artist	Candidate days	Suppressed	Accepted	Complete windows
BTS	42	24	18	18
BLACKPINK	47	32	15	14
TWICE	30	15	15	15
Stray Kids	22	10	12	12

Artist	Candidate days	Suppressed	Accepted	Complete windows
SEVENTEEN	35	21	14	14
aespa	41	23	18	18
Total	217	125	92	91

Warm-up dates are not eligible for detection. Complete windows contain day 0 plus 28 later days. Within-artist primary follow-up windows do not overlap.

3 2 Confirmed recovery durations

The pooled Kaplan–Meier median confirmed return is day 8. Among the 91 complete episodes, 42 confirm return by day 7 (46.2%), 64 by day 14 (70.3%) and 78 by day 28 (85.7%). The complete-episode recovery classes therefore comprise 42 brief episodes, 22 intermediate episodes, 14 prolonged episodes and 13 unresolved episodes. Intermediate denotes confirmation on days 8–14 and prolonged denotes confirmation on days 15–28. The unresolved class records failure to observe confirmation within the horizon, without asserting that those episodes never return.

The curve declines sharply in the first week and then more slowly. At day 3, the first confirmations can occur by construction. This measurement delay matters when interpreting very short responses: a three-day confirmation cannot report a one-day recovery even if the day after onset already falls near the baseline. The median return is an operational duration under this confirmation rule, rather than a physiological or psychological attention span. Risk-set counts displayed with the curve include episodes censored or confirmed at each elapsed date and make the denominator change explicit.

The artist descriptions show heterogeneity within the selected panel. BTS confirms 16 of 18 complete episodes by day 14, whereas Stray Kids confirms 7 of 12 and aespa 11 of 18. The artist-specific median confirmed return ranges from day 5 for SEVENTEEN to day 12 for Stray Kids and aespa. These samples are small and differ in episode mix, baseline conditions and calendar timing. The results describe the observed pages; they cannot establish that one artist’s audience intrinsically remains attentive longer than another’s. The mean of the six artist-specific day-14 fractions is 69.6%, close to the pooled 70.3%, so equal artist weighting does not reverse the broad result.




Figure 4. Time to confirmed return under a frozen reference. The product-limit curves use all accepted episodes and explicit censoring. Numbers in the artist legend are episode counts. Risk sets are shown immediately before the specified elapsed day. Curves do not carry independent-episode confidence bands because temporal dependence and small artist cohorts would make that presentation misleading.

Table 3. Artist recovery and residual burden

Artist	N	Day 14 return	Day 28 return	Median burden	Median tail share
BTS	18	88.9%	94.4%	8.47	23.5%
BLACKPINK	14	64.3%	78.6%	6.84	36.6%
TWICE	15	73.3%	93.3%	9.63	40.4%
Stray Kids	12	58.3%	75.0%	10.81	51.5%
SEVENTEEN	14	71.4%	92.9%	9.16	42.1%
aespa	18	61.1%	77.8%	13.73	58.7%

N is the number of complete episodes. Return fractions use complete episodes. Burden is positive excess in baseline equivalent days over days 0–28. Tail share uses positive excess on days 8–28. Artist-level values are descriptive.

3.3 Residual attention and peak intensity

The median early peak is 2.79 times the frozen baseline. The median normalized positive excess burden is 9.80 baseline equivalent days and its mean is 13.75. The difference between median and mean reflects a right-skewed burden distribution, including several persistent episodes. The equal-artist mean burden is 13.67, again close to the episode-weighted mean. These quantities measure accumulated deviations rather than continuous duration, so the numerical resemblance between a burden and a recovery day must not be interpreted as identity between them.

The median tail share is 37.5% and the mean is 39.4%. Thirty-one of the 91 complete episodes accrue more than half of their positive excess after day 7. Since the tail contains 21 dates and the initial period eight, the

median tail-to-initial daily excess ratio is also calculated and equals 22.8%. Thus, the typical tail's daily positive deviation is much smaller than its initial counterpart, while its longer observation period still contributes appreciably to total burden. The trajectory summary combines declining intensity with continued accumulated attention, rather than labelling every nonzero tail as a maintained peak.

Figure 5 preserves the signed normalized trajectories. There are 702 below-baseline observations among complete episode trajectories and 51 observations above the selected early-peak excess. Thirteen episodes have a later count exceeding their first-week maximum. Those later observations explain why a universally monotone decay description would be inappropriate. The shaded interquartile range shows variation across episodes and is not a confidence interval for the median. The separate artist medians and their equal-weight average provide a check on the pooled profile's dependence on artist composition.




Figure 5. Normalized signed excess after onset. Normalization divides $Y - B$ by the maximum first-week excess $P - B$. Curves use the 91 complete episodes. Negative values and later values greater than one are retained. The shaded region is the episode interquartile range, and the dashed line averages the six artist-specific medians.




Figure 6. Complete episode recovery classes by artist. Brief means confirmation by day 7, intermediate days 8–14, prolonged days 15–28 and unresolved no confirmation by day 28. Each horizontal bar is normalized within the artist's complete cohort. Small denominators limit comparison between artists.

Early-peak ratio and normalized burden have a descriptive Spearman rank correlation of 0.335 across complete episodes. This indicates imperfect agreement between intensity ordering and accumulated-excess ordering within the panel. Both measures contain the early peak and share a baseline denominator, so this correlation is not an independent predictive validation. No significance claim or causal interpretation is attached to it. Figure 7 instead supplies concrete examples of why a high peak and a long residual burden should remain separate fields in a monitoring record.

The largest early-peak ratio occurs for BTS on 12 June 2024. Its first-week maximum is 21.06 times the reference, yet return is confirmed on day 6 and the tail share is 3.2%. The BTS onset of 4 March 2026 has a smaller early peak of 5.50 times the reference, but its burden reaches 79.80 baseline equivalent days and 89.3% accrues in the tail. No frozen-reference return is observed by day 28. These are observed dates and traffic contrasts. This analysis does not infer their causes from the artist name or match them speculatively to releases or news.




Figure 7. Early intensity and accumulated positive excess. Each point is a complete episode. Both axes use logarithmic scales. Labels identify selected observed contrasts. Peak and burden share mathematical components, so their association is a descriptive comparison rather than independent predictive evidence.

3 4 Reference adaptation changes recovery classifications

The adaptive comparison confirms 52 of 91 complete episodes by day 7, 75 by day 14 and 90 by day 28. The corresponding frozen-reference counts are 42, 64 and 78. At day 14, 64 episodes are confirmed under both references, 11 only under the adaptive reference, none only under the frozen reference and 16 under neither. At day 28, 78 are confirmed under both, 12 only under the adaptive reference and one under neither. The net adaptive-minus-frozen day-14 difference is 12.1 percentage points.

The 60-day calendar-block interval for that paired difference is 5.5–19.0 percentage points. Resampling therefore preserves a positive difference within this historical set under the primary block specification. The difference is a change in classification produced by the chosen reference, not an estimate of false recovery against a labelled truth. Updated medians may reflect a genuine shift in ordinary page traffic as well as absorption of elevated observations. Retaining both references makes the definition visible to an analyst who needs to distinguish return to the old level from normalization relative to recent conditions.

This paired design matters because changing both onset detection and return definition at once could confound the comparison. Here the onset list, observation window and required number of low days are identical. Only the reference in the confirmation condition changes. The result does not imply that an adaptive detector should be abandoned. It identifies the question answered by each part of the workflow: adaptation can be useful when recognizing a current departure, while a frozen reference is useful when tracking the subsequent deviation from the alert's original context.




Figure 8. Frozen and adaptive reference comparisons. Left: cumulative confirmed-return curves for the same 92 accepted episodes with censoring. Right: complete-cohort fractions and counts at days 7, 14 and 28. A daily updated reference changes the return condition but does not alter the onset list.

Table 4. Paired complete episode return counts

Horizon	Both references	Adaptive only	Frozen only	Neither
Day 7	42	10	0	39
Day 14	64	11	0	16
Day 28	78	12	0	1

Each row covers the same 91 complete episodes. Adaptive and frozen references use the identical three-day confirmation length. The difference concerns measurement definitions rather than an independently verified true recovery state.

3.5 Peak decline and renewed attention

Ninety complete episodes confirm a half-excess decline after the selected first-week peak. Among those confirmed declines, the median lag after the peak is four days. This conditional median excludes the one episode without an observed half-excess confirmation and uses a different time origin from the baseline-return curve. It should not be directly read as an uncensored median recovery duration. Twelve episodes confirm the partial decline but still do not confirm frozen-baseline return by day 28. Figure 12 presents the two clocks separately.

The persistent BTS episode beginning on 4 March 2026 illustrates that distinction especially clearly. Its early peak is on day 0, and it confirms a half-excess decline on day 3. Nevertheless, it remains unconfirmed against the original return tolerance throughout the four-week window. A large fall from the peak can therefore coexist with substantial residual traffic. Monitoring based only on peak reduction would discard information about the size and timing of the remaining deviations.

Seventy-three complete episodes have a return by day 21 and therefore a full seven-day renewal opportunity. Fourteen of these episodes, or 19.2%, contain a later count above 1.5 times the onset reference in that opportunity window. The other 18 complete episodes are not renewal negatives: 13 lack a confirmed

return and five confirm too late to observe seven additional dates. The renewal flag consequently applies to a defined subgroup after return and cannot be interpreted as the fraction of all surges that eventually recur.

Artist-specific renewal counts range from one of nine eligible BLACKPINK episodes to five of 17 eligible BTS episodes. This small-denominator result is reported as counts as well as proportions. The interval between two surges in real news attention is not estimated, because the flag uses one fixed seven-day opportunity and a count threshold. Renewed traffic may arise from related or unrelated activity. The daily counts alone cannot establish whether two peaks concern the same promotional episode or the same readers.




Figure 9. Renewed attention after confirmed return. Labels show renewal counts divided by eligible complete episodes. Eligibility requires a return by day 21 so all seven subsequent days are observed inside the day-28 horizon. Renewal means a later daily count above 1.5B, without applying the original robust-score condition again.

Table 5. Opportunity adjusted renewed attention

Artist	Complete	Eligible after return	Renewed	Renewal fraction
BTS	18	17	5	29.4%
BLACKPINK	14	9	1	11.1%
TWICE	15	14	1	7.1%
Stray Kids	12	8	2	25.0%
SEVENTEEN	14	11	3	27.3%
aespa	18	14	2	14.3%
Total	91	73	14	19.2%

The denominator includes only episodes with seven fully observed days after confirmation. Unresolved and late-confirming episodes do not become zero-valued renewal outcomes.

3.6 Sensitivity and calendar block intervals

Detector changes alter the episode catalogue substantially. Lowering the onset ratio to 1.5 produces 140 episodes, while raising it to 2.5 produces 61. The 56-day baseline yields 91 episodes and a score threshold of four yields 87. Either tested scale-floor change leaves the 92 primary onsets unchanged. These alternatives demonstrate dependence on operational definitions rather than reveal an empirically optimal setting. The original threshold remains the primary specification throughout the manuscript.

Changes to the return definition affect the result even without changing the onsets. Tightening the tolerance from 1.25B to 1.10B reduces the day-14 return fraction from 70.3% to 46.2%, while loosening it to 1.50B increases it to 84.6%. Requiring only one low day produces a day-14 fraction of 78.0%. It also permits a much earlier confirmation and is more exposed to a temporary dip. The choice of tolerance and confirmation length therefore expresses the monitoring question; it cannot be presented as a universal property of attention.

Shortening spacing to 14 days yields 104 accepted episodes with overlapping follow-up, while lengthening it to 43 days yields 83. Their day-14 complete-episode return fractions are 68.9% and 67.1%. Restricting the main catalogue to 2024–2026 onsets gives 56 episodes, 55 complete, with 76.4% confirmed by day 14 and 87.3% by day 28. This later-period summary preserves the historical detector state and is not an independent external replication. The differences indicate changes in the selected period and episode mix, which also qualify the approximate stationarity assumption underlying block resampling.




Figure 10. Sensitivity to detector and return definitions. Filled points show complete-episode return by day 14 and hollow points day 28. The primary rule is highlighted. Every non-period alternative changes one factor. The full numerical table additionally includes the two scale-floor changes, which leave the primary results unchanged. Fourteen-day spacing permits overlapping windows.

Table 6. Primary and alternative operational rules

Specification	Episodes	Complete	Day 14 return	Day 28 return
Primary	92	91	70.3%	85.7%
2024 to 2026 onsets	56	55	76.4%	87.3%
Baseline 56 days	91	90	72.2%	84.4%
Onset ratio 1.5	140	139	76.3%	87.8%
Onset ratio 2.5	61	60	68.3%	83.3%
Robust score 4	87	86	68.6%	84.9%
Scale floor 0.05	92	91	70.3%	85.7%
Scale floor 0.20	92	91	70.3%	85.7%
Return ratio 1.10	92	91	46.2%	70.3%
Return ratio 1.50	92	91	84.6%	93.4%
One day confirmation	92	91	78.0%	90.1%
Refractory 14 days	104	103	68.9%	83.5%
Refractory 43 days	83	82	67.1%	84.1%

The later-period row restricts primary onsets rather than redetecting a shortened panel. Return-rule changes keep primary onsets unchanged. The 14-day spacing alternative permits overlapping follow-up windows.

The 60-day calendar-block interval for the complete-episode day-14 return fraction is 59.6–81.2%. The day-28 interval is 76.7–93.4%. The mean burden interval is 10.77–16.66 baseline equivalent days, and the mean tail-share interval is 34.1–44.5%. All 2,000 replicates in each of the three block-length specifications retain at least one event and are valid. Table 7 reports the primary intervals and the sensitivity of the paired reference difference to block length. This uncertainty describes resampling of the observed episode records; it does not add measurement truth that the public count data lack.

Table 7. Calendar block resampling intervals

Quantity	Point estimate	95% lower	95% upper
Return by day 14	70.3%	59.6%	81.2%
Return by day 28	85.7%	76.7%	93.4%
Mean burden in baseline days	13.75	10.77	16.66
Mean tail share	39.4%	34.1%	44.5%
Adaptive minus frozen at day 14	12.1%	5.5%	19.0%
Paired difference with 30 day blocks	12.1%	5.8%	19.4%
Paired difference with 90 day blocks	12.1%	5.8%	18.8%

Primary intervals use 60-day circular calendar blocks, 2,000 replicates and seed 20261003. Percent differences are percentage points. Calendar dates are resampled jointly across artists with their real episode outcomes attached. Intervals are descriptive and conditional on the selected historical panel.

3.7 Observed contrasting trajectories

Three episodes are selected by recorded numerical rules to make the contrasts inspectable. The brief-return example is the shortest confirmed duration, breaking ties by greatest normalized burden. It is SEVENTEEN on 15 August 2026, whose onset peak reaches 15.80 times the baseline. Return confirms on day 3, yet the accumulated burden is 16.83 baseline equivalent days. A large single-date excess can thus produce substantial accumulated burden even when the recovery clock is short. Its tail share is only 8.7% and it has no seven-day renewal flag.

The persistent example is the complete unresolved episode with greatest normalized burden, BTS on 4 March 2026. The renewed example is the eligible episode with the greatest subsequent seven-day ratio, SEVENTEEN on 23 June 2024. In the latter, return confirms on day 3, but a later first-week peak on day 5 reaches 6.16 times the original reference. The sequence shows why confirmation must remain a time-stamped state rather than a retrospective assertion that the entire later trajectory is low. Figure 11 includes the adaptive median alongside the observed counts so the reference's evolution can be examined directly.

These examples are chosen to explain distinct outcomes and are not representative or randomly sampled. Their peak dates and quantities are available in the full catalogue, allowing the selection rules to be reproduced. No external event labels are added to the figure. This keeps the empirical contrasts separate from speculation about which release, appearance or news story generated attention.




Figure 11. Observed examples of brief persistent and renewed attention. Each panel shows visits divided by the frozen onset median B . The dashed line is the updated median divided by B , not the adaptive return threshold. The dotted horizontal line is the fixed $1.25B$ return tolerance. Vertical lines mark confirmed frozen-reference return where observed. Cases are selected by predetermined descriptive rules.

Table 8. Illustrative episodes with different recovery trajectories

Description	Artist	Onset	Peak ratio	Return day	Burden	Tail
Brief return	SEVENTEEN	2026-08-15	15.80	3	16.83	8.7%
Persistent excess	BTS	2026-03-04	5.50	>28 censored	79.80	89.3%
Renewed attention	SEVENTEEN	2024-06-23	6.16	3	11.26	11.1%

Onsets are UTC dates. Peak ratio uses the largest count on days 0–6. The censored entry records that confirmation is unobserved by day 28, without estimating an exact later duration. Case selection is specified in Methods.




Figure 12. Partial peak decline and confirmed baseline return. Left: the 78 complete episodes with observed baseline return and half-excess decline, using two distinct time origins. Right: counts with each confirmation by day 28 among all 91 complete episodes. Half-excess confirmation occurs after the selected early peak and is not an exponential half-life.

4 Discussion

4.1 What the separated measurements reveal

The principal finding is that the status of an attention episode depends on which aspect of the trajectory is being measured. The main catalogue shows substantial early recovery, with roughly seven in ten complete episodes confirming return within two weeks. It also shows a meaningful residual component: about one third place more than half their positive excess after the first eight days, and a minority remains unresolved at the four-week horizon. Those statements coexist because a group can have a falling daily deviation while still accumulating a material total over many later dates.

The contrast between peak decline and baseline return gives the workflow practical value. A high peak creates a permissive half-excess threshold, so reaching half the excess can be relatively easy while the series remains far above ordinary pre-alert traffic. The twelve complete episodes that satisfy the partial criterion without satisfying full return make this issue empirically visible. A monitor that records only a percentage drop from the peak would lack the historical level needed to assess the remaining deviation.

Conversely, a single very large but brief increase can have high total burden without demanding a long recovery duration.

Reference adaptation produces a second measurement distinction. The increase from 78 to 90 day-28 confirmations arises from comparing traffic with a progressively updated recent level. This is useful information rather than an automatic indictment of adaptation. A stable new level can be ordinary in recent conditions while still differing from the alert's original context. The analyst must decide which question is operationally relevant and name it accordingly. The paired design demonstrates the size of that choice without mixing it with a different set of detected events.

The renewal result adds a third temporal state. Fourteen of the 73 episodes with enough later observation cross the renewal threshold after confirmation. That does not invalidate their earlier confirmed return. It shows that the confirmation describes a completed sequence available on that date, rather than the end of all subsequent attention. Retaining the return date and the subsequent flag creates a record that can be updated without rewriting its past state. This would be useful for a dashboard that needs to distinguish a new increase from an alert that has never returned.

4 2 Implications for audience monitoring

A practical monitoring view should retain the alert date, frozen reference, current ratio, confirmed-return status, accumulated positive excess and the latest renewal opportunity. At onset, it can show the robust score and count ratio that justified the alert. During follow-up, it can show whether enough consecutive low dates have completed, while avoiding a future-known peak or an eventual duration label. After confirmation, it can begin a fixed seven-day check and explicitly mark whether that opportunity is complete. A last-day alert should remain pending rather than being converted into a negative outcome by the end of a spreadsheet.

The frozen reference and adaptive reference can be shown side by side with their definitions. An unresolved frozen comparison and a confirmed adaptive comparison then becomes an interpretable situation: traffic is recently ordinary but still has not approached the old reference under the primary rule. This prevents a change of denominator from being silently read as audience disengagement. The burden field separately helps quantify historical volume above that reference. Its unit remains baseline equivalent days, so it should not be relabelled as new fans, campaign lift or revenue.

The locally retained research package makes the measurement workflow inspectable by the author and publisher. It combines public data identity, explicit information timing, a reproducible detector, administrative censoring, documented denominators, numerical figure tables and an editable research workbook. Its single-episode explorer supports changes to the return tolerance and inspection of the resulting confirmation date, while the main study catalogue remains tied to its fixed primary protocol. Every summary can be traced locally to an episode record and then to its observed daily counts, supporting inspection of unusual trajectories beyond the aggregate plots. The working package is retained privately rather than deposited with the public article.

4 3 Limitations and interpretation boundaries

The six pages are a selected convenience panel, and their visitors are not a representative sample of Korean music fans. English Wikipedia is one information channel with its own reading habits, article identity and platform changes. Daily user-classified visits can include repeated visits and cannot resolve people,

geography, intent or sentiment. The panel does not include individual artists' other pages, translations, song pages or all related redirects. Generalization beyond these specific pages requires a different sampling and page-identity design.

The alert rule is operational. Without independently annotated events, this study cannot estimate false positives or missed real-world surges. The 29-day spacing intentionally suppresses closely spaced candidates, which means prolonged and renewed increases can remain within a single episode. The sensitivity alternatives expose part of this dependence but do not validate an optimal threshold. Calendar coincidence with a release or announcement would also be insufficient to infer cause. A causal study would need verified event information and a defensible comparison design rather than merely a peak in the same week.

The fixed median is a historical reference and is not the unobserved traffic that would have occurred without an event. Broader shifts, secular trends or a change in ordinary interest can leave the series elevated against that reference even after a particular stimulus has faded. The adaptive comparison clarifies this ambiguity rather than resolving it. Positive excess truncation also has a specific consequence: later below-baseline counts do not offset earlier positive ones. This suits a positive-burden description but cannot be substituted for a net causal effect.

Return timing depends on the three-day confirmation and 25% tolerance. The earliest duration is fixed at three days by construction, and the sensitivity table shows that a narrower tolerance materially changes the fraction labelled recovered. The 28-day horizon leaves 13 complete episodes unresolved and cannot characterize their eventual tails. The single panel-end episode is additionally censored after two days. Complete-window burden summaries therefore apply to the 91 eligible episodes, while the product-limit curve retains all 92. These are distinct analysis populations and their denominators are kept visible.

The renewal flag is restricted to a seven-day post-confirmation opportunity and a threshold of 1.5B. It does not establish a new qualifying primary alert or whether the same underlying event has returned. Artist-level denominators are particularly small. Likewise, the early-peak normalized trajectory is retrospective and may be exceeded by later peaks. The Spearman association between peak and burden shares mathematical inputs and is not evidence that a predictive model has been validated. Exponential or power-law decay parameters would be poorly justified for many of the observed nonmonotone paths.

Calendar-block intervals preserve some observed dependence and shared-date structure, but they condition on the detected historical catalogue. They neither incorporate uncertainty about a different alert definition nor overcome changes in the distribution of attention across years. The 30- and 90-day results provide a block-length check, not proof that all relevant dependence has been captured. The data were previously available for related work, so the analysis makes no claim of an untouched holdout test. These limits qualify the conclusions while leaving the measured within-panel contrasts reproducible.

5 Conclusion

A past-only surge rule applied to 10,404 daily artist-page visits identifies 92 separated episodes, including 91 with complete four-week observation. The primary frozen-reference median confirmed return is day 8, and 85.7% of complete episodes confirm by day 28. Residual attention remains material for some episodes, with a median tail share of 37.5%. A partial decline after an early peak can occur without return to the previous level, and a confirmed return can precede renewed attention. Updating the reference raises the

day-28 confirmation count from 78 to 90 on the same complete episodes, making the definition of recovery consequential.

The study’s contribution is a reproducible distinction between an alert, its intensity, its confirmed recovery and its later positive burden. It supports transparent page-traffic monitoring with explicit observation opportunities and sensitivity checks. The findings remain descriptive of selected English Wikipedia pages and do not identify unique fans, causes or promotional returns. Keeping the baseline and the information clock visible is essential when interpreting how attention subsides.

Data Code and Author Statements

Frankie Gao is the sole named author. The affiliation is perspectiveopen.com and the correspondence address is zcfbfga@ucl.ac.uk. OpenAI Codex assisted with analysis code, numerical computation, visualizations and manuscript drafting. The sole named author is responsible for reviewing and approving the manuscript and its interpretations. All reported computations were executed locally on preserved public numerical data. No empirical value is supplied for an unobserved follow-up date.

The numerical source data are publicly available through the Wikimedia Analytics API under CC0 1.0 [2], using the artist-page titles, dates and request settings documented in this paper. The locally executed research package contains the unchanged source CSV and 12 raw API responses, original retrieval provenance and hashes, the analysis protocol, complete code, executed notebook, episode and trajectory tables, sensitivity results, figures and an editable spreadsheet. These working materials are retained privately by the author and publisher and are not deposited as public supplementary files. The package contains no private participant records. The analysis uses aggregate public traffic and makes no individual-level inference. Readers may direct methodological and reproducibility enquiries to the corresponding author.

References

- [1] Wikimedia Foundation. Page view analytics. Wikimedia Analytics API documentation. Accessed 3 October 2026. <https://doc.wikimedia.org/generated-data-platform/aqs/analytics-api/reference/page-views.html>
- [2] Wikimedia Foundation. Access policy and data licensing. Wikimedia Analytics API documentation. Accessed 3 October 2026. <https://doc.wikimedia.org/generated-data-platform/aqs/analytics-api/documentation/access-policy.html>
- [3] Crane R, Sornette D. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences*. 2008;105(41):15649–15653. doi:10.1073/pnas.0803685105. <https://pmc.ncbi.nlm.nih.gov/articles/PMC2572957/>
- [4] Lorenz-Spreen P, Mønsted BM, Hövel P, Lehmann S. Accelerating dynamics of collective attention. *Nature Communications*. 2019;10:1759. doi:10.1038/s41467-019-09311-w. <https://www.nature.com/articles/s41467-019-09311-w>
- [5] National Institute of Standards and Technology. Median absolute deviation. *Dataplot reference manual*. Updated 11 April 2016. Accessed 3 October 2026. <https://www.itl.nist.gov/div898/software/dataplot/refman2/auxillar/mad.htm>
- [6] Kaplan EL, Meier P. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*. 1958;53(282):457–481. doi:10.1080/01621459.1958.10501452. <https://www.tandfonline.com/doi/abs/10.1080/01621459.1958.10501452>
- [7] Künsch HR. The jackknife and the bootstrap for general stationary observations. *The Annals of Statistics*. 1989;17(3):1217–1241. doi:10.1214/aos/1176347265. Author publication record verified 3 October 2026. <https://people.math.ethz.ch/~hkuensch/papers/>