



DEMOCRACY
REPORTING
INTERNATIONAL

DISRUPT
DISINFO
NETWORKS



Detecting Inauthentic Coordination and Foreign Pipelines on Telegram:

What We Found During Hungary's 2026 Elections



About Democracy Reporting International

DRI is an independent organisation dedicated to promoting democracy worldwide. We believe that people are active participants in public life, not subjects of their governments. Our work centres on analysis, reporting, and capacity-building. For this, we are guided by the democratic and human rights obligations enshrined in international law.

DRI's Digital Democracy Programme Unit is a research-based team conducting practical analysis of online political discourse and threats to its plurality, authenticity, factfulness, and free participation in these spaces.

Acknowledgements

This report was written by Democracy Reporting International. Our analysis benefitted from the input and contextual and technical expertise of Konrad Bleyer-Simon and Salvatore Romano. This report is part of the [Disrupt Disinfo Networks](#) project, funded by the German Federal Foreign Office. Its contents do not necessarily represent the position of the German Foreign Federal Office.

Julieta Jimenez designed the layout of this publication.

Date: June 2026



This publication is available under a Creative Commons Attribution Non-Commercial 4.0 International license.

Julieta Melina Jiménez designed the layout of this publication.



Table of Contents

Executive Summary	4
Key findings	5
What's uncommon in @common_sens?	6
A case study in multilingual coordinated inauthentic behaviour	
Four dimensions of coordinated inauthentic behaviour	7
Single origin	7
Operational automation and synchronisation	9
Misrepresentation of identity	12
Artificial amplification	12
Mapping the Telegram eco-system	16
Balanced search, skewed findings	16
Domain analysis: Content pipelines and manufactured reach	19
Network analysis: How does the network proliferate content?	21
Narrative analysis: extraction and diffusion	25
Narrative extraction	26
Cross-platform diffusion	28
Implications and recommendations	31
Telegram's accountability gap under the Digital Services Act	31
Recommendations	32
Methodology	34





Executive Summary

The role the Telegram messaging app plays as a hub for foreign information manipulation and interference (FIMI) operations and coordinated disinformation campaigns is increasingly under scrutiny. The platform's combination of low moderation, public group infrastructure, and a lack of substantial regulation makes it especially suitable as an organising space for extremist groups, conspiracy communities, and foreign actors seeking to influence political debates. Countering these networks requires understanding their prevalence and impact across online spaces, including how actors use Telegram to coordinate and promote certain narratives. The risks these networks pose is increasingly heightened around electoral periods, and detection has remained an ongoing challenge.

The 2026 Hungarian parliamentary elections provided a relevant case study for investigating these dynamics. During our monitoring, we discovered a coordinated inauthentic network of 19 Telegram channels masquerading as independent news outlets while publishing identical content simultaneously across 19 languages, including Hungarian. Our wider investigation of the Telegram ecosystem surrounding the elections was no less significant; despite using ideologically balanced keywords to identify channels, our analysis revealed a platform environment structurally skewed toward far-right and Russian-aligned actors, foreign media pipelines, and coordinated amplification patterns across 77 channels and reaching over 1.4 million members, collectively.

These findings contribute to a wider discussion on whether Telegram should fall within the enhanced obligations applicable to Very Large Online Platforms (VLOPs) under the EU's Digital Services Act. At present, the platform remains exempt, and the identified networks operate in the gap between what existing monitoring can detect and what existing law can regulate.



Key findings

- **Networks without networking:** We discovered 19 seemingly unrelated channels that, as later examination proved, shared the same infrastructure, published the same content in different languages, misrepresented themselves as news outlets, and exploited automation scripts to publish content simultaneously.
- **Structural asymmetry in the identified Telegram ecosystem:** We identified 77 public Telegram channels actively sharing content about the 2026 Hungarian elections through a balanced keyword-based approach. Despite the ideologically balanced keyword selection, the resulting Telegram ecosystem was disproportionately dominated by far-right, conspiratorial, and Russian-aligned channels, suggesting an asymmetric visibility within the observed election-related environment.
- **Content pipelines and manufactured reach:** The analysis of the shared URLs revealed amplification strategies between Telegram channels sharing the same content within seconds, as well as a documented content pipeline from Kremlin media to the identified Hungarian information space on Telegram.
- **Convergence of structural centrality and content redistribution:** The Telegram network surrounding the 2026 Hungarian elections was organised around a small number of highly connected channels, with one channel simultaneously acting as a major cross-community bridge and a primary distributor of translated foreign content, Russian military analysis, and sanctioned media sources.
- **Many narrative formulations, one coherent image:** From the full corpus of Telegram messages relevant for the Hungarian elections, we extracted 292 macro-narratives combining between two and 17 micronarratives and messages formulations each. From a qualitative assessment of the macro-narratives, Hungary is depicted as a sovereign actor resisting pressure from European institutions and globalist networks, particularly in relation to the war in the Ukraine and economic matters.
- **Cross-platform narrative analysis revealed presence, but not directionality:** While different narrative formulations extracted from Telegram were observable across multiple online environments, no consistent evidence emerged regarding where narratives systematically originated or appeared first.



What's uncommon in @common_sens? A case study in multilingual coordinated inauthentic behaviour

During manual data collection for the monitoring of the Hungarian elections, a researcher noticed the same message appearing simultaneously across multiple channels in different languages. This was only possible to spot because the messages contained names of key political figures – (incumbent Prime Minister Viktor) *Orbán*, (opposition party Tisza leader Péter) *Magyar*, (ruling party) *Fidesz* – terms that do not translate and, therefore, appear identically, regardless of the surrounding language. Had we relied solely on algorithmic collection at scale, this cross-language repetition would have gone entirely undetected. What began as a side observation turned into a dedicated investigation.

We uncovered a network consisting of 19 Telegram channels posting identical content simultaneously across 19 languages, each presenting itself as an independent local news source. We initially analysed seven directly observed channels. A keyword-based search subsequently confirmed that the network extended to at least 19 channels, independently verified by an external expert. Every indicator of coordinated inauthentic behaviour identified in the initial seven channels held true across all 19. If the network proves larger still, we have no reason to expect different results.

By standard logic of network analysis, these 19 channels would not register as a network at all. They don't link to each other, and they don't forward each other's content at any meaningful rate. Measured by forwarding relationships – the usual method of mapping connections on Telegram – they appear as 19 unrelated channels. What actually connects them is invisible using that approach – identical content, identical timing, identical visual identity, and a shared administrative infrastructure. This section documents those connections, proposes a framework for detecting similar networks, and demonstrates how that framework can be applied and confirmed quantitatively. Future researchers encountering similar patterns will, therefore, have a replicable procedure to follow.

These findings also speak directly to the broader research question motivating this report. Narrative dissemination research typically asks how content moves across platforms and between actors. This network inverts that question. It demonstrates how the same narrative can be made



to appear as if it is moving organically across independent communities, when it is actually being distributed from a single point, simultaneously. Understanding that mechanism is as important for narrative dissemination research as understanding genuine diffusion because the two are designed to be indistinguishable from the outside.

Four dimensions of coordinated inauthentic behaviour

We've organised the indicators in this section around four dimensions of coordinated inauthentic behaviour: single origin, operational automation and synchronisation, misrepresentation of identity, and artificial amplification. No single indicator is conclusive on its own. A shared visual identity could be a coincidence. Coordinated posting schedules could reflect editorial practice. Purchased subscribers are not uncommon among legitimate channels seeking rapid growth. What makes the combination of indicators documented here amount to coordinated inauthentic behaviour is their accumulation; each additional indicator that holds true reduces the probability that the pattern is organic, and the combined weight of evidence ultimately leaves no other reasonable explanation. We've grouped the indicators below by dimension, noting how each was assessed and what it would look like in a network without coordination.

Single origin

Any coordinated inauthentic network operating across multiple apparent identities requires, by definition, a single actor or coordinated group behind those identities. Establishing a single origin is therefore the foundational dimension of any CIB assessment. Four indicators speak to this dimension in the Common_Sense network, each assessable through different methods and each reinforcing the others.

The most immediately visible indicator is **how consistent the channel names are across languages**. Each of the 19 channels carries a name that expresses the same concept in its respective language: @common_sens_news, in English; @die_stimme_der_vernunft in German; @józan_ész, in Hungarian; @robuste_bon_sens, in French; @_mintede_țăran, in Romanian; @köylü_mantığı, in Turkish; @akiliya_kutumia, in Swahili; and equivalents across the remaining 12 languages. The names are not direct translations of each other, but occupy the same semantic space, evoking practical wisdom or grassroots good sense. This consistency across 19 languages is a pattern that requires a deliberate naming decision applied uniformly when the channels were created.

Visual identity leads to the same conclusion. All initially observed channel logos followed an identical design pattern: channel names written across two rows, with black as the primary colour and red as the secondary. Several variants incorporate curly brackets as a shared design element. The logos are visually distinct enough to appear independent at first glance, yet similar

enough to be immediately recognisable as a family when viewed together. This design consistency held across 18 of the 19 channels, with the Bulgarian channel @zdrav_razum_BG being the only visually distinct variant.



Figure 1. Channel logos

The shared contact mechanism is the most operationally significant indicator of single origin. Seventeen of the 19 channels list the same Telegram bot, @Olej_w_Glowie_bot, as their contact point. The remaining two channels, @selsky_rozum and @púchoviny, share a second bot, @selsky_rozum_bot. Channels presenting themselves as independent local news sources in different countries don't share contact bots. This points directly to a single administrative structure.

Network growth patterns provide a fourth distinct line of evidence. Creation dates were obtained for 14 of the 19 channels through Telemetrio – a Telegram analytics platform used to track channel growth, engagement, advertising performance, and audience trends – with the remaining five estimated through fixed-interval systematic sampling. The data shows that 12 channels were created in two distinct waves, with six channels appearing on the same day in each wave. Channels created independently by unrelated actors in different countries do not cluster their creation dates in this way.




Figure 2. Network growth over time

Single origin alone doesn't prove coordinated inauthentic behaviour. Media organisations operating local editions in multiple languages share infrastructure, contact mechanisms, and visual identity, by design. What distinguishes legitimate multilingual operations from coordinated inauthentic ones isn't simply a shared origin, but what that the shared origin is used for – whether the channels are transparent about their relationship, whether their growth reflects genuine audience interest, and whether their content and publishing patterns reflect independent editorial judgment. The indicators documented in the following dimensions – automated synchronised posting, artificial amplification, and systematic misrepresentation of identity – are what transform shared origin from an organisational feature into a CIB indicator.

Operational automation and synchronisation

Nineteen channels publishing identical content across 19 languages cannot be operated manually at any meaningful scale. The volume of content alone makes human coordination implausible; between 1 January and 10 April 2026, the 19 channels published a combined 222,154 messages. Two quantitative assessments confirm both the automation and synchronisation of the network.

We assessed automation by examining the **intervals between consecutive messages** within each channel. Automated posting produces intervals between messages clustered around fixed values, rather than varying freely as they would under manual editorial scheduling. The figure below shows the interval distribution across all 19 channels. Intervals cluster sharply at multiples of 300 seconds: five minutes, ten minutes, fifteen minutes, and so on, up to one hour. The spike at zero seconds likely reflects messages containing multiple media files recorded as separate entries. This pattern, consistent across all 19 channels, is the signature of a scheduling system publishing at fixed intervals, rather than a human editor making individual posting decisions.




Figure 3. Message interval distribution

We assessed synchronisation through two complementary approaches. The first examined what **proportion of channels posted within the same minute** at any given point in the observation period. Under the hypothesis of synchronisation, the expectation is that when one channel posts, all or nearly all others will post simultaneously. The figure below shows the result; publishing events cluster at near-100-per-cent channel participation. Across thousands of publishing events throughout the observation period, almost every channel posted within the same minute as every other.




Figure 4. Channel participation per publishing event

The second synchronisation assessment used the **Jaccard index**, a standard measure of set similarity ranging from 0 to 1, where 0 indicates no shared elements and 1 indicates identical sets. Applied here, the Jaccard index measures the proportion of minutes in which any two channels posted simultaneously relative to the total minutes when either posted. For channels publishing at the observed frequency of approximately 12,000 messages each over a 100-day window (144,000 possible minutes for publication), two independently operating channels would share posting minutes purely by chance in roughly 4 to 5 per cent of cases, yielding a Jaccard index of approximately 0.04.¹ The values we observed across channel pairs in this network range from 0.5 to 0.93 (50 per cent to 93 per cent similarity).




Figure 5. Pairwise Jaccard index values

One limitation of timestamp-based synchronisation assessment deserves mention; if two channels publish different content within the same minute by coincidence, this analysis would register it as coordination. In a dataset with no prior evidence of simultaneous cross-language publications of identical content, this limitation would matter significantly. Here, where such a posting pattern had already been established through direct observation, the risk is acceptable. More importantly, the Jaccard index values observed are not marginally above the chance threshold, but rather exceed it by a factor of 15 to 20. Such a gap cannot be explained by coincidence.



Misrepresentation of identity

Coordination and automation describe how this network operates. Misrepresentation of identity is what makes it inauthentic. The network presents each of its 19 channels as an independent, locally rooted news source in its respective language. That presentation is false in every verifiable dimension.

The clearest evidence is the most straightforward; a basic search confirms that **none of the 19 channels exists outside Telegram** – no website, no registered media organisation, no editorial history, no journalistic track record. They publish no bylines, carry no ownership disclosure, and provide no information about who produces their content. The watermark on the English-language channel reads @common_sens_news, explicitly positioning it as a news outlet. Channels presenting themselves as news sources while having none of the infrastructure that news sources carry are lying to their audiences about what they are.

The channel-naming convention compounds the misrepresentation. Names evoking practical wisdom or grassroots common sense in each language are designed to suggest locally rooted, independently minded sources. A reader encountering @józan_ész, in Hungarian, @minte_de_țăran, in Romanian, or @kansan_maalaisjarki, in Finnish, has no reason to connect those channels to each other or to suspect they share an origin. The names are chosen to appear local. That appearance is how the coordination remains invisible to its target audience.

The Telemetrio platform, which provides audience analytics for Telegram channels, tags the majority of channels in this network as cheaters, due to fake views. Telemetrio's methodology for this assessment is not publicly documented, and this finding is, therefore, presented as supplementary rather than primary evidence. Its relevance lies in that a third-party service, operating independently for commercial rather than research purposes, reached the same conclusion about the inauthenticity of this network's audience metrics.

Misrepresentation of identity most directly answers the question of why this network constitutes coordinated inauthentic behaviour, rather than simply coordinated behaviour. Coordination is not inherently problematic. News organisations coordinate. Political parties coordinate. What distinguishes this network is that the coordination is actively concealed behind the appearance of independent local journalism. That concealment is the mechanism through which the network's content acquires the legitimacy it would not have if its true nature were visible.

Artificial amplification

The three dimensions documented so far establish that this network was built by a single actor, operates automatically and synchronously, and misrepresents itself as independent local journalism. This dimension addresses whether the network's apparent reach reflects genuine audience interest.



At the time of observation, the 19 channels had a combined subscriber count of 167,059.

Channel handle	Language	Subs on Apr 5
@stimme_der_vernunft	German	21,401
@selskyrozum	Czech	13,360
@olej_w_glowie	Polish	9,482
@common_sens_news	English	8,464
@los_pies_en_la_tierra	Spanish	8,147
@seljacka_logika	Serbian	8,125
@jozan_esz	Hungarian	7,972
@koinos_nous	Greek	7,944
@robuste_bon_sens	French	7,918
@sale_in_zucca	Italian	7,873
@minte_de_taran	Romanian	7,776
@puchoviny	Slovakian	7,720
@gek_genoeg	Dutch	7,488
@tvunat_ha_kfar	Hebrew	7,414
@common_sens_arab	Arabic	7,396
@koynu_mantigi	Turkish	7,375
@kansan_maalaisjarki	Finnish	7,363
@zdrav_razum_bg	Bulgarian	7,164
@akili_ya_kutumia	Swahili	6,677
TOTAL		167,059

Table 1. Language and number of subscribers per channel.

The German-language channel @stimme_der_vernunft alone accounted for 21,401 of those subscribers, by far the largest individual channel in the network. Its growth data explains why. Between 25 September and 28 September 2025, the channel grew from 1,707 to 42,633 subscribers, a gain of over 40,000 in three days. It subsequently lost subscribers, settling at 21,401 by the time we observed it. **Bulk subscriber purchase** is the only explanation consistent with that trajectory.




Figure 6. Telemetrio growth rates screenshot for the German-language channel @stimme_der_vernunft

The remaining channels show similar if less dramatic patterns, with sharp subscriber spikes concentrated in February 2026 across all channels available on Telemetrio.

Artificial amplification, like misrepresentation of identity, matters because of what it produces in the audience. A channel with 40,000 subscribers carries more perceived credibility than one with 1,700, regardless of how those subscribers were acquired. In a network whose entire mechanism depends on appearing independently legitimate across multiple communities, inflating the subscriber counts can also be seen as a deliberate manipulation strategy.



Dimension	Indicator	Method/Measure
Single origin	Consistent channel names	Direct observation/translation
	Consistent visual identity	Visual confirmation
	Shared contact mechanism	Manual channel exploration/ Telemetrio
	Network growth patterns	Telemetrio data/fixed-interval systemic sampling in scraping
Automation and synchronisation	Shared publishing intervals	Inter-message interval distribution
	Synchronised event participation (all)	Network-normalised publishing events, participation in events (each minute where at least one channel published = event)
	Synchronised event participation (between channels)	Pairwise Jaccard index
Misrepresentation of identity	No existence outside of Telegram	Google search
	Channel names semantics	Interpretation of translation
	Fake views	Telemetrio assessment/ method undisclosed
Artificial amplification	Bulk purchase of subscribers	Telemetrio information
	Shared growth trajectories of channels	Telemetrio information

Table 2. Dimensions, indicators, and the methods of their assessment used in the research.



Mapping the Telegram eco-system

Our analysis of the Telegram information space surrounding the 2026 Hungarian elections draws on **77 public channels** identified through a combination of keyword-based search and forwarding relationship mapping. These channels collectively **reached 1,391,429 non-unique members** at the time of observation.

Balanced search, skewed findings

Channel identification was based on a list of 116 Hungarian-language keywords (see the full list of keywords in the Appendix) developed with the support of a local digital research expert, and organised into three groups:

- **Election-related keywords:** neutral references to the electoral event itself, such as “national elections”, “campaign period”, and “electoral system”;
- **Topic-based keywords:** associated with key issues of the electoral period, including “immigration”, “Ukraine War”, and “EU budget veto”; and/or
- **Ideological keywords:** terms capturing more explicitly politically aligned discussions. To ensure balance, we included an equal number of ideological keywords to reflect arguments across all sides of the political spectrum – government, opposition, conservative, and progressive.

Each keyword was queried through Telegram’s public channel search, with the first ten channels being noted down. Later, the channels were retained for analysis if they contained five or more different keywords from the list, regardless of the keyword’s type.

The results reflected an uneven distribution. Despite the deliberate ideological symmetry of the keyword set, the **first 31 channels** identified were predominantly associated with far-right movements and Russian-aligned media. Ideological keywords occurred the most frequently across all monitored channels, accounting for 47 per cent of all keyword occurrences, followed by topic-related keywords, at 34 per cent, and election-related keywords, at 18.6 per cent. Among the top-ranked channels by keyword co-occurrence, the first two displayed affiliations with QAnon and presented themselves as alternatives to mainstream media, as shown in Table 2.

 A NAGY ÉBREDÉS 6 386 subscribers A kritikai gondolkodást a RENDSZER (politika, közgazdaság, pénz, pszichológia) működési szintjén kell megalapozni. Míg ezt nem értjük meg, és a problémákat a JÓ és ROSSZ ellentétpárra építjük és ennek örökös harcával magyarázzuk a világ nem fog megváltozni	Keyword co-occurrence: 73 (1st in the list) Channel Title (translated): The Great Awakening Channel scription: “Critical thinking must be grounded in the operating level of the SYSTEM (politics, economics, money, psychology). Until we understand this and build problems on the pair of opposites of GOOD and EVIL and explain them with the eternal struggle, the world will not change”.
--	---

Table 3. Screenshot of the top identified channel by keyword occurrence.

The fourth-ranked channel described itself as a media outlet covering events “from Russia and the wider world in the light of the Special Military Operation”.

Several channels carried an official or semi-official character, including:

- ➔ The far-right Our Homeland (Mi Hazánk) party, represented both through its official channel and through what appears to be the official outlet of its news platform, Magyar Jelen;
- ➔ The far-right Vadhajtasok blog, closely aligned with the governing party, including both its official channel and the personal profile of its founder, Bede Zsolt;
- ➔ The official channel of the pro-government media outlet Pesti Sracok;
- ➔ The channel of Hungarian blogger Tamas B. Nagy, who is frequently cited in far-right media; and
- ➔ The official channel of the Russian Embassy in Hungary.

Figure 7 shows the top 15 seed channels by keyword occurrence. The channels ranking the highest, @ebredes2017, @greatawakeningmagyarok, @Zsuzsa2021, and @oroszkazigazsagoldalan, are among those identified as affiliated with QAnon, Russian-aligned media, or far-right movements, reinforcing the pattern described above.




Figure 7. Top 15 seed channels (of 31) by keyword occurrences

To extend the range of actors under observation, we analysed all forwarded messages within the election-relevant corpus, and extracted the original source channels. This produced **46 additional channels**, referred to throughout as linked channels, which were included in subsequent analysis. Among these, the channels with the highest number of total forwards were associated with political figures from the far-right Our Homeland party, followed by conspiratorial and anti-globalist groups.

Seed and linked channels together reached a total audience of **1,391,429 non-unique members**, with linked channels accounting for 88.8 per cent of that total, as shown in Figure 8.




Figure 8. Top 15 channels by size in the final dataset

These numbers alone do not constitute evidence of coordination, but they are consistent with a pattern in which the channels actively redistributing content within this network are drawing from sources with substantially larger audiences than their own.

Domain analysis: Content pipelines and manufactured reach

Domain analysis examines the external links shared across the monitored channels, revealing both where content originates before entering the network and how channels in this ecosystem manage their presence on other platforms.

The most shared domain in the entire dataset is vadhajtasok.hu, a site identified by our local expert as a pro-Fidesz outlet with far-right editorial positioning, accounting for 2,842 shares across eight channels. While its dominance alone is notable, what makes it analytically significant is the behaviour of two channels in relation to it, namely **@Zsuzsa2021** and **@ebredes2017**. These channels repeatedly posted identical URLs from vadhajtasok.hu within seven to 11 seconds of each other across multiple dates in October 2025. On 13 October, both channels posted the same article at 17:48, precisely five seconds apart and approximately four hours after the original post appeared on vadhajtasok.hu itself. This pattern, where two channels wait and then post simultaneously hours after the source, **suggests the channels are being operated by a single actor manually promoting content rapidly.**

The single URL generating the most views in the entire dataset leads to a **profile page on max.ru**, a Russian state-backed platform developed by Russian social media platform VKontakte.

That profile belongs to Dmitry Smirnov, a Kremlin pool correspondent for the pro-Putin newspaper *Komsomolskaya Pravda*² and operator of a large pro-Kremlin Russian language Telegram channel @PUL N3 (“@Пул N3”).³ A single link to this profile generated 7.1 million views and 22,901 forwards, approximately three times more than any other URL in the dataset. The explanation lies in how the content was distributed: @PUL N3 content, consisting of pro-Orbán statements and anti-EU framing **originally produced in Russian, was being translated into Hungarian** and redistributed through @oroszokazigazsagoldalan, @Zsuzsa2021, and @ebredes2017, with every redistributed message ending with an invitation to subscribe to the max.ru profile. This represents a documented content pipeline from Russian state media into the Hungarian information space, with identifiable channels serving as translators and distributors.

A separate and related finding concerns **rybar.ru**, a Russian military analysis outlet with a Hungary-specific channel, @Rybar_HU, which produces localised content, including high-quality infographics. The channel @oroszokazigazsagoldalan, which, as noted above, serves as a distributor of Russian state media content, also functions as the primary amplifier of Rybar content, cross-posting it within seconds of the original publication. In one documented case the gap between the original @Rybar_HU post and the @oroszokazigazsagoldalan repost was six seconds. The same channel also shared links to arabic.rt.com, the Arabic-language service of RT, through obfuscated URLs consistent with an attempt to avoid detection as sanctioned content. That @oroszokazigazsagoldalan performs all three of these functions simultaneously, distributing Russian state media, amplifying Russian military analysis, and sharing sanctioned RT content, makes it one of the most analytically significant channels in the entire dataset, and connects directly to the network finding that it is also the primary bridge between otherwise separate communities within the Telegram network, as will be shown in the next section.

The Russian Embassy in Hungary maintained a documented presence across Facebook, VKontakte, and X simultaneously, with its social media handles appearing together in the same messages and generating identical metrics across all three platforms: 33 shares, 37,931 views, and 88 forwards. Those posts were amplified by @oroszokazigazsagoldalan and @vdtahu, among others. Separately, a Hungarian political blogger operating the channel @magyarbtamas was soliciting donations through boosty.to, a Russian crowdfunding platform that allows creators to receive payments outside the sanctions regime. The channel was amplified by @oroszokazigazsagoldalan and @hirekpentrado. At least one donor reported their bank flagging the transfer as suspicious.

The **cross-platform behaviour** documented in the domain analysis extends beyond Russian-aligned content. MyLordBebo, an English-language far-right commentator with close to 800,000 followers on X, was systematically routing his own X content into Telegram using **fixvx.com**, a mirror service that converts X posts into embeddable links, bypassing X’s login wall. That content, covering pro-Viktor Orbán, anti-Ukraine, and anti-EU themes, was then redistributed into Hungarian channels by @oroszokazigazsagoldalan, @Zsuzsa2021, and @ebredes2017. The channel operated continuously from October 2025 through polling day on 12 April, including live coverage of the election result. In the opposite direction, the X profile of @InfoDefMagyarok was the single most shared URL in the entire dataset, generating 664,627 views from 1,025 shares



across five channels, despite the account having only 68 followers and having been inactive for nearly a month at the time of analysis. Its view count on Telegram exists entirely because channels within this network were pushing traffic toward it. A similar pattern was observed with the X account of **@geopolitics_prime**, which generated 722,307 views from only 60 shares across two channels. Both cases point to the same phenomenon – external platform presence being artificially inflated from within Telegram, rather than reflecting genuine audiences elsewhere.

Finally, **the most shared Facebook link** in the dataset leads to a page presenting itself as an NGO called **Hungarian Legion**, which appears on inspection to be the neo-Nazi paramilitary group “Legio Hungaria”, an organisation with documented links to far-right extremist groups across Europe.⁴ The same page dominates the most shared Instagram and X links in the dataset, with identical share counts of 27 across four channels. The exact numbers of appearances of one account across platforms should be understood as the consequence of all platforms being simultaneously shared in the same message, and not much more than that.

Network analysis: How does the network proliferate content?

The network analysed consists of 71 Telegram channels connected through 1,424 forwarding interactions, drawn from a corpus of election-relevant messages of which 4.7 per cent contained forwarded content. Connections between channels (edges) are directional, running from the channel doing the forwarding toward the channel whose content was forwarded, and weighted by how many times that forwarding occurred. One important property of this network is that every channel that forwarded content from others was itself forwarded at least once, meaning no channel in the dataset functions as a purely one-way distributor. This should be interpreted carefully; it reflects the structure of what the dataset captured, rather than a universal property of the broader Telegram ecosystem.

The network as a whole is sparsely connected, with only three per cent of all possible connections between channels being active. This means forwarding behaviour is selective; channels do not forward broadly across the network, but concentrate their activity on specific other channels. At the same time, the network is not fragmented. All 71 channels form a single connected structure, with no isolated actors present. Within that connected structure, **five distinct communities** emerged via clusterisation (highlighted in different colors in the graph); groups of channels that forward among themselves more than they forward across group boundaries. This combination – sparse overall connectivity with strong internal clustering – suggests that the network is organised around distinct but internally cohesive groups, rather than a single unified flow of content.




Figure 9. Network graph. Node size represents the number of forwards received; edges are directed from the channel doing the forwarding toward the channel whose content was forwarded; colors represent the five identified clusters

Within that structure, the channels attracting the most forwarding attention are the Our Homeland Movement's leader, László Toroczka, and its news platform Magyar Jelen, both terminal destinations of amplification, rather than active redistributors. The channel doing the most connective work across communities, @oroszkazigazsagoldalan (roughly translated to *Russians on the side of truth*), is also the channel documented in the domain analysis as a primary distributor of translated foreign content, a pipeline for Russian military analysis, and a conduit for sanctioned media. That convergence of structural centrality and external content distribution in a single channel is the most analytically significant finding of this section.

To understand how content circulates within this network, it helps to think of channels as playing distinct roles. Some channels are primarily destinations – their content gets forwarded by others, but they rarely forward content themselves. Others are primarily distributors – they actively forward content from elsewhere, but generate little forwarding attention in return. A smaller



group does both, receiving significant forwarding attention, while also actively redistributing content across the network. And the majority of channels do neither at any meaningful scale, participating in the network, but without significant influence in either direction. These four roles are formalised in the analysis as content hubs, aggregators, active hubs, and marginal channels, respectively, and they structure the findings that follow.

Role	Definition	Count	Share of network
Marginal	Channels with limited forwarding activity in both directions. Present in the network, but without significant influence on content flow.	43	60.6%
Aggregator	Channels that forward content from others extensively, but receive little forwarding attention themselves. Active distributors, rather than content destinations.	11	15.5%
Content hub	Channels whose content is widely forwarded by others, but which rarely forward content themselves. Destinations, rather than distributors.	10	14.1%
Active hub	Channels that both forward extensively and receive significant forwarding attention. The primary cross-community redistributors in the network.	7	9.8%

Table 4. Full role distribution.

The table below shows the ten most intense forwarding relationships in the network, measured by the total number of forwards between pairs of channels. The most active relationship – by a considerable margin – is **@mihazank forwarding to @torockzai**, with 216 forwards recorded. @mihazank is also the most extreme aggregator in the dataset – a channel that made 377 total forwards toward other channels while receiving comparatively little forwarding attention in return. @torockzai, by contrast, is a content hub, as the most forwarded channel in the entire network with 325 total forwards received, belonging to Our Homeland Movement leader Torockzai.



Source	Target	Weight
@mihazank	@toroczkai	216
@szentkoronaradio_official	@varmegye	82
@matrixhungary	@magyarjelen	71
@mihazank	@novakelod	52
@mihazank	@lantosjanos	49
@zsuzsa2021	@magyarjelen	47
@ebredes2017	@magyarjelen	47
@oroszokazigazsagoldalan	@rybar_hu	46
@hirekpentrado	@toroczkai	42
@greatawakeningmagyarok	@toroczkai	40

Table 5. Highest forwarding relationships.

As anticipated, content attention in this network is highly concentrated. A small number of channels attract the vast majority of forwarding activity, while most receive very little, visible in the gap between the average number of forwards received across all channels (20.6) and the median (5). @toroczkai and @magyarjelen, the movement's news platform Magyar Jelen, are the two most amplified channels by a significant margin, leading both in total forwards received (weighted in-degree) and in overall network influence (PageRank).⁵ The Russian Embassy channel in Hungary receives fewer forwards than either of them, but the channels forwarding it are among the most active forwarders in the network, giving it a PageRank score that its raw forward count alone would not suggest. The figure below illustrates this distribution; the blue bar shows total forwards received for each of the top 20 channels, while the orange bar shows how many distinct channels forwarded them. The gap between the two bars, most visible for @toroczkai, shows that high amplification volume in this network tends to come from a small number of dedicated sources, rather than broad support.




Figure 10. Top 20 channels by weighted in-degree

The channels that received the most forwards are not, however, the ones doing the most connective work within the network. @toroczkai and @magyarjelen, despite their dominance by volume, have a betweenness score of zero, meaning they sit at the ends of forwarding chains, rather than along the paths connecting different parts of the network. Betweenness measures how often a channel lies on the shortest path between two other channels; a high score means that removing that channel would disconnect parts of the network that currently communicate through it. The two channels with the highest betweenness scores are @oroszokazigazsagoldalan and @hirekpentrado, neither of which ranks among the top channels by forwarding volume. Their significance lies not in how much content they receive, but in where they sit – between communities that would otherwise have limited connection to each other. @oroszokazigazsagoldalan, in particular, combines this bridging role with the content pipeline functions documented in the domain analysis, making it the single most analytically significant channel in the dataset across both dimensions of analysis.

The following section turns from the structure of this network to the content circulating within it, examining what narratives were present, how persistent they were, and whether they can be traced across platforms beyond Telegram.

Narrative analysis: extraction and diffusion

From the full corpus of Hungarian election-relevant posts, the analysis identified 292 political macro-narratives that groups individual messages conveying the same underlying meaning. These macro-narratives serve as a primary unit of analysis throughout this section. Time distribution analysis showed that some narratives remained persistent over the months approaching

the elections (particularly those centred on criticism of the EU), while others appeared only episodically, emerging in response to specific events and disappearing once those events lost salience. To investigate how narratives appeared across platforms, we selected 28 macro-narratives following two criteria: 1) narratives having more than one formulation or 2) a higher risk of information politicisation or manipulation related to the described topic or event. The analysis also differentiated between the two mechanisms of cross-posting, where the same actor publishes identical or near-identical content on different platforms simultaneously, and organic diffusion, where different actors independently share the same narrative. This distinction is important for understanding how narratives spread, although, as the analysis will show, the available data allows only a partial answer to that question. Thus, our investigation also advances our methodological understanding⁶ of how cross-platform narrative diffusion can be studied empirically.

Narrative extraction

Before conducting any cross-platform comparison, the analysis first identified and described the narratives emerging within the Telegram corpus itself. This section maps what narratives were present in the dataset, what they contained, and how they were distributed over time.

Starting from the full corpus of Telegram messages related to the Hungarian elections, messages were clustered semantically to group together different formulations conveying the same underlying meaning. The analysis identified the highest-engagement messages within each cluster, and used them to characterise the corresponding macro-narrative. This process resulted in 292 macro-narratives in total, each represented by a varying number of message formulations. The table below indicates what kind of different messages can be semantically traced to belong to the same narrative.

Example of narrative and its formulations

Macro-narrative	Formulation 1	Formulation 2	Formulation 3
Criticism of European Union policies	“Orbán sharing the insight: European leaders are morons, he says, in elegant language.”	[extract] “The war that we Europeans have not been able to stop for four years will soon be ended by the Americans. Europe would continue and even expand the war. This step is equivalent to an open declaration of war.”	[extract] “Orbán warned that Brussels pushes forced diversity, yet ignores the value of a cohesive, secure community, stressing that Hungary will not surrender to population replacement.”

Table 6. Example of three formulations of the same macro-narrative.

Our analysis shows that the narrative corpus turned out to be ideologically consistent. The messages frequently portrayed Hungary as a sovereign actor resisting pressure from European institutions, Ukraine, or “globalist networks”, particularly on questions of foreign politics and international conflict, energy independence, and national identity. The domestic political dimension centres on the conflict between the government and the opposition, led by Péter Magyar, including references to scandals such as the mishandling of personal data by his Tisza party, which ultimately won the elections, through its app, reportedly exposing the data of up to 200,000 Hungarian citizens. Some narratives adopted strongly polarising terms and contain claims that are difficult to verify. Examples include allegations that Ukrainian president Volodymyr Zelenskyy was transferring USD five million weekly to fund the Tisza party, or that Orbán received payments from Russian President Vladimir Putin in exchange for political alignment. The corpus also included narratives around real events that appeared susceptible of instrumentalisation, such as reports claiming that intercepted Ukrainian armoured convoys were transporting cash and gold to interfere in the elections.

The corpus divides into two temporal patterns, visible in the figure below. Some narratives maintain continuous presence across the entire observation period; EU criticism, controversial domestic policies, and Hungary-Russia diplomatic relations appear at varying intensity across every month in the dataset. Others are strictly episodic, spiking in response to specific events, and then disappearing. Topics such as EU sanctions against Russia and Ukrainian military activity follow this pattern.

Both types were concentrated in March, the month before the elections, with episodic narratives peaking on top of a persistent ideological frame that had been present throughout the observation period, as presented in Figure 11, below.




Figure 11. Top 15 narratives distribution over time (monthly based)

Cross-platform diffusion

To verify whether the narratives identified in the Telegram corpus appeared beyond Telegram, we used Google Search indexing as a practical proxy for cross-platform presence. Verbatim queries based on different formulations of 28 selected macro-narratives, chosen for both the number of available formulations and their potential information politicisation and manipulation risk, were used to retrieve the top ten results per query across platforms and domains. This approach has a specific practical value beyond platform-native search; it approximates what an ordinary person would find if they searched for a story they had encountered, and it captures content hosted on proprietary websites that platform searches would miss.

Nevertheless, Facebook emerged as the dominant platform in those results, consistent with the platform's prominence in Hungary. Beyond Facebook, narratives appeared on other social media platforms such as X, Instagram, YouTube, and TikTok with a single occurrence on the Russian social media platform VKontakte. Within news media and outlets, results included *magyarjelen.hu*, *Kuruc.info*, *hungary.news-pravda.com*, and *24.hu*. The blog category included *magyarbtamas.com*, *hazairend.hu*, and *polkorrekt.com*. The distribution of results across platforms is shown in the figure below.




Figure 12. Distribution of Google SERP results by platform

Where the same handle appeared across both platforms, we classified the occurrence as cross-posting by the same actor. Where handles differed, we treated it as potentially involving different actors, with the caveat that the same political actor may operate different accounts across platforms, making this distinction indicative, rather than definitive. The figure below shows the distribution of narratives by diffusion type across the top 15 of the 28 selected narratives.




Figure 13. Top 15 macro-narratives by diffusion type

To complement this comparison, the analysis considered aggregated trend data from Meta for a selection of episodic narratives, allowing their timing on Facebook to be compared against their first appearances in the Telegram dataset. Three specific examples of contents were examined in this way.




Figure 14. Telegram and Facebook activity, Meta censorship narrative

For the narrative concerning alleged censorship by Meta (Figure 14), an initial instance appeared on Telegram in mid-February, coinciding with marginal activity on Facebook. The peak of discussion on both platforms occurred from 11 March onward, coinciding with the suspension of the Facebook page Crime Hunters. This is the closest the data comes to a traceable triggering event shared across platforms.

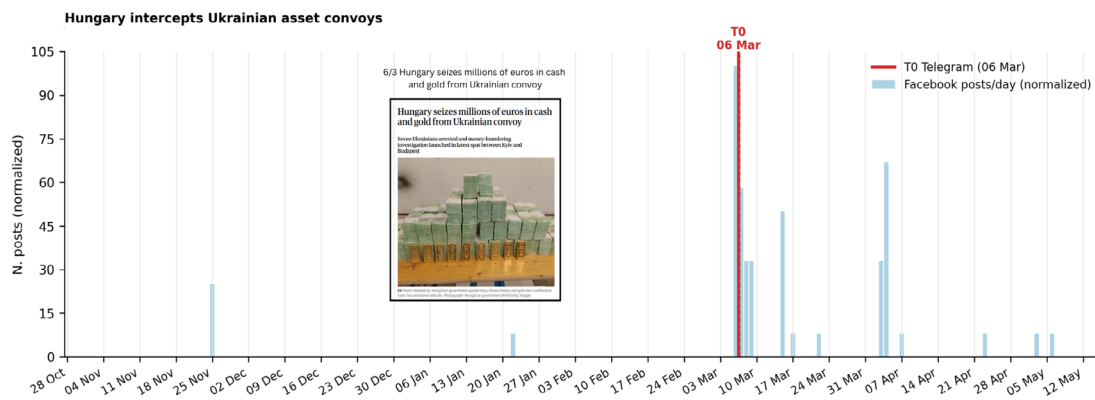


Figure 15. Telegram and Facebook activity, Ukrainian convoy narrative

The narrative about Ukrainian armoured convoys transporting cash and gold (Figure 15) shows a spike on both platforms around 6 March, the date of the event. No sustained activity followed on either platform, confirming this as a purely episodic narrative, tied to a single news moment.

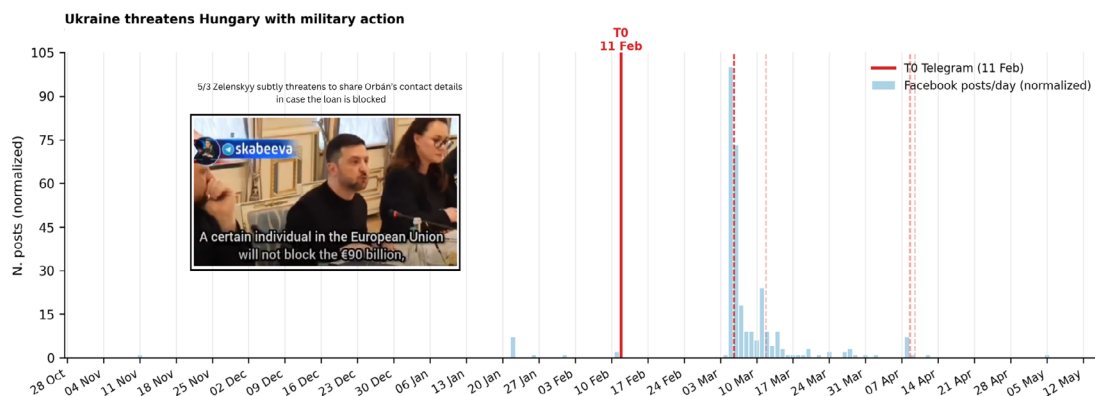


Figure 16. Telegram and Facebook activity, Zelenskyy threats narrative

The narrative concerning threats by Zelenskyy toward Orbán shows no clear pattern across the observation period, with activity scattered across both platforms without any consistent temporal relationship between them.

Across all three cases presented here, and consistent with the broader timestamp comparison, no directional pattern emerged. Narratives appeared on both platforms simultaneously, rather than sequentially. This means the data neither confirms nor rules out Telegram as an origin point; the temporal resolution available does not allow that conclusion to be drawn. What the cross-platform analysis can say with confidence is more modest; the narratives carrying the highest assessed manipulation risk were present and findable across multiple platforms, their appearance was broadly simultaneous, and whether that simultaneity reflects coordinated cross-posting, independent parallel circulation, or simply the speed at which content moves once published cannot be determined with the available data.

Implications and recommendations

Telegram's accountability gap under the Digital Services Act

Our findings add to a growing body of evidence of literature documenting influence operations and coordinated inauthentic behavior on messaging platforms, and underscore Telegram's relevance to the information environment surrounding a high-stakes democratic event – the 2026 Hungarian parliamentary elections⁷.

Although Telegram channels fall within the scope of the DSA, the platform does not meet the 45 million monthly active user threshold required for designation as a VLOP. This means it is not subject to the DSA's most stringent obligations, including conducting systemic risk assessments and implementing mitigation measures. This regulatory gap is compounded by Telegram's historically minimalist approach to content moderation, which offers users and affected communities little recourse when "harmful but legal" content or – such as in the case we found – inauthentic behaviour circulates at scale.

That said, the European Commission has confirmed it is scrutinising Telegram's methodology for counting users, and independent research and Members of the European Parliament have questioned the accuracy of Telegram's self-reported figures, suggesting the platform may already exceed the threshold.⁸

This results in a platform with considerable reach and potential influence in European Union operating under a significantly lighter regulatory burden than its reach, influence, and risk would warrant.

The majority of our findings do not constitute illegal content. We believe, however, that they **pose a systemic risk to the integrity of the Hungarian electoral process and civic discourse within the meaning of Article 34(1)(c) of the DSA**. Had equivalent findings emerged on a VLOP or VLOSE, they would have triggered obligations to act, and likely enforcement scrutiny from the European Commission.

- ➔ We found strong evidence of coordinated inauthentic behaviour on Telegram. The inauthentic use of services through fake accounts, bots, deceptive use of a service, and other automated or partially automated behaviours are explicitly identified by the DSA as examples of systemic

risk under Article 34(1).⁹ Telegram's own verification framework¹⁰ for public figures and organisations requires channels to demonstrate a verifiable external presence, yet all 19 channels in the network we found, which present themselves as news sources, have no presence outside of Telegram, no editorial accountability, and no track record as legitimate media outlets.

- We identified at least one channel actively bypassing EU sanctions against Russian-associated media outlets, operating without apparent restriction on the platform. Article 2f of [Council Regulation 833/2014](#) explicitly prohibits operators from transmitting or distributing content from sanctioned outlets through any means, including online platforms. RT Arabic is listed as a sanctioned outlet under [Council Regulation 2023/427](#). Telegram's failure to restrict access to RT Arabic and allowing users within the EU to access this content, therefore, constitutes a direct violation of that prohibition.

This qualifies as illegal content under the DSA and gives rise to a systemic risk within the meaning of Article 34(1)(a). Following a civil society report published in 2025, Telegram blocked access within the EU to several Russian state-owned media channels.¹¹ The continued presence of the RT Arabic channel reflects a gap in Telegram's moderation approach and warrants urgent attention by the Belgian Digital Services Coordinator.

- We found channels associated with neo-Nazi movements and groups showing conspiratorial features, propaganda, and material that may contain information manipulation, including claims Zelenskyy was channelling €5 million per week to fund the opposition Tisza party, and that Ukrainian armoured convoys had been intercepted carrying cash and gold to finance illegal interference in Hungarian politics. Although thorough content analysis falls outside the scope of our research, such illustrative examples may indicate a systemic risk to electoral integrity under the DSA Electoral Integrity Guidelines¹². Recently, the European Board for Digital Services also identified the spread of such extremist narratives and harmful conspiracy theories as prominent and recurrent systemic risks.¹³

Recommendations

To Telegram

- **Block and maintain restrictions on channels broadcasting sanctioned Russian outlets within the EU**, including language-specific or regional variants of sanctioned outlets, which our findings suggest are used deliberately to circumvent existing blocks;
- **Sign and implement the Code of Conduct on Disinformation**. Telegram's accession would bring it within a structured accountability framework and signal a meaningful commitment to addressing disinformation at scale, particularly during electoral periods;



→ **Strengthen content-moderation policies in line with the DSA Electoral Integrity Guidelines.** Specifically, this could include:

- policies to detect and address coordinated inauthentic behaviour, including behavioural signals, such as anomalous subscriber growth patterns consistent with bulk purchases; and
- channel verification requirements for channels presenting themselves as news sources, applied proactively rather than only on application, drawing on Telegram's own existing verification framework, which already requires channels to demonstrate an external presence and press coverage.

To the Belgian Institute for Postal Services and Telecommunications, as the competent Digital Services Coordinator for Telegram

- **Require Telegram to comprehensively identify and maintain restrictions on channels broadcasting sanctioned Russian outlets within the EU.** Broadcasting sanctioned outlets constitutes illegal content under EU law.

To the European Commission

- **Reassess Telegram channels' regulatory status under the DSA, and evaluate whether it meets the threshold criteria warranting designation as a VLOP;**
- **Update the DSA Electoral Integrity guidelines to include suggested mitigation measures for non-VLOPs that pose demonstrable systemic risks to electoral integrity and civic discourse.** The current framework operates on the premise that size is the exclusive proxy for risk, yet our findings show that a non-VLOP platform can materially affect electoral integrity; and
- **Require VLOPs and VLOSEs to establish structured cooperation with Telegram and other relevant platforms to address cross-platform disinformation,** including participation in incident-sharing mechanisms and election integrity readiness dialogues.

To VLOPs and VLOSEs

- **Actively participate in cross-platform cooperation mechanisms with Telegram and other non-VLOPs, as explicitly called for by the DSA, the Electoral Integrity Guidelines, and the Code of Practice on Disinformation.** In practice, this means sharing intelligence on cross-platform influence operations, coordinating on incident response during electoral periods, and proactively flagging the migration of malicious actors to smaller platforms.



Methodology

In this section, we outline the methodological steps we used to conduct the analyses presented in this report, complementing and expanding what has been described in the main text.

Workflow 1: Seed and linked channel identification and messages data collection

To build the dataset of public Telegram messages, we began by identifying the set of public channels relevant to our research. We used a keyword-based approach through the search functionality available in Telegram Premium (see the Annex for the full 116 keyword list). The keyword list was developed in collaboration with a local digital expert involved in the project.

After compiling an initial list of channels, we applied a co-occurrence criterion, retaining only channels matching at least five keywords. This threshold was introduced to increase the likelihood that the selected channels were substantively related to the research topic. In addition, two channels previously identified through parallel investigations were included, despite not meeting the filtering criteria. This process resulted in a final set of 31 seed channels.

Using the Telegram API, we extracted all publicly available content published by these channels during the six months preceding the election date. Non-relevant posts were then filtered out through a mixed-method approach that included an LLM-assisted classification step, manual validation, and reliability testing. From the resulting set of Hungary election-related posts (32,094), we extracted 46 unique linked channels identified through the `forwarding_from` metadata field, when available.

We subsequently repeated the same collection and filtering procedure for the linked channels. Finally, we merged all Hungary election-related messages published by both seed and linked channels between 1 October 2025 and 15 April 2026, resulting in a final dataset of 37,400 messages.

Workflow 2: Domain analysis

The analysis of external links was designed to examine how Telegram channels connect to wider online information environments and which external sources are most recurrently embedded within election-related conversations. Rather than treating URLs as simple references, the analysis

approached them as indicators of informational alignment, amplification behaviour, and cross-platform sourcing practices within the observed Telegram ecosystem.

To support this analysis, all outbound URLs contained in the Hungary election-related Telegram corpus were extracted and consolidated at the domain level.

The most recurrent domains were subsequently enriched with contextual, cross-platform, and infrastructural information intended to support analytical assessment, potentially under a FIMI-oriented perspective. A qualitative assessment was conducted on a subset of highly recurrent domains and URLs, in order to evaluate the functional role they occupy within the broader Telegram information environment.

Workflow 3: Network analysis

This step aimed to understand the relationship between seed and linked channels, and to analyse the election-related content redistribution within the identified Telegram network. Particularly, in this weighted network, nodes represented Telegram channels, while directed edges represented forwarding relationships between channels. Edges' weights were calculated based on the frequency of forwarding interactions observed during the data-collection period.

The resulting network consists of 71 channels connected through 1,424 forwarding interactions. Only messages containing observable forwarding metadata were included in the construction of the network. As a consequence, the network captures redistribution patterns visible within the collected dataset, rather than the entirety of information flows occurring across Telegram.

To assess the structure of the network, weighted in-degree and weighted out-degree were used to measure forwarding attention and forwarding activity, respectively. PageRank was used to identify structurally prominent channels within the forwarding network, while betweenness centrality was used to identify channels positioned along paths connecting otherwise weakly connected parts of the network.

Community detection was performed on the weighted directed network using the Louvain algorithm to identify groups of channels characterised by denser internal forwarding interactions relative to external connections.

Workflow 4: Topic and narrative extraction

To identify recurring topics and narratives within the corpus of Hungary election-related Telegram messages collected from both seed and linked channels, we adopted a multi-stage analytical pipeline combining unsupervised topic modelling with LLM-assisted semantic analysis.

In this experiment, the LLM used was the Mistral AI model medium.

Following an initial preprocessing stage, aimed at reducing noise and standardising message content, semantic clustering was performed using multilingual MiniLM BERT embeddings. This process generated 100 initial topic clusters, including a residual noisy cluster (-1) containing messages that could not be reliably assigned to a coherent semantic group.

To prioritise content with higher visibility and circulation potential within the Telegram ecosystem, an engagement score was computed for each message, based on a weighted combination of interaction metrics, including views, forwards, and replies. Using these scores, a fixed number (50) of highly engaged messages was sampled from the noisy cluster, while proportional sampling was applied across all remaining clusters. This procedure produced a final subset of approximately 2,000 top-engaged messages for deeper analysis.

Each selected message was then processed using an LLM-assisted extraction step designed to identify micro-narratives, defined as concise semantic representations of the specific event, claim, or interpretative framing conveyed within an individual message. To facilitate systematic comparison and aggregation, micro-narratives were then normalised into a structured actor/action/target format. This standardisation step enabled the transformation of heterogeneous message content into a more analytically consistent structure suitable for iterative clustering and higher-level narrative aggregation.

The resulting structured micro-narratives were subsequently re-clustered and semantically consolidated through an additional LLM-assisted aggregation phase, producing a second analytical layer of macro-narratives. This generated a hierarchical interpretative structure linking individual messages to micro-narratives and, subsequently, to broader narrative formations.

Workflow 5: Cross-platform narrative travel analysis

From the corpus of macro-narratives extracted from the Telegram dataset, we identified 28 narratives combining both a high number of associated formulations (micro-narratives and messages) and a comparatively higher potential for information manipulation or political instrumentalisation. These narratives were selected for a cross-platform diffusion analysis (the full list is available in the Annex of this report).

To investigate how these narratives circulated beyond Telegram, individual micro-narrative formulations were used as verbatim search queries within Google Search Engine Results Pages (SERPs). Using SERP-based retrieval enabled a broader observation of narrative presence across both social media platforms and independently hosted websites, including news outlets and blogs. This approach was intended to capture narrative replication and cross-platform circulation patterns, rather than platform-specific dynamics. By tracing the appearance of semantically consistent narrative formulations across indexed web sources, the analysis aimed to identify



whether and how Telegram narratives were also identified into broader online information environments. Several limitations should, nevertheless, be acknowledged. First, temporal reconstruction remains imperfect, as Google Search does not consistently provide reliable indexing timestamps and, in many cases, only displays publication dates, whose accuracy and comparability may vary across sources. As a consequence, the approach is better suited to identifying cross-platform narrative presence than to reconstructing precise diffusion timelines. Second, while the method enables the detection of narrative occurrences across platforms, it does not provide reliable estimates of narrative volume or reach within individual platforms. The analysis records the presence of a narrative instance, rather than the scale of its circulation, engagement, or audience exposure within the hosting environment.

To partially address these limitations, the SERP-based analysis was complemented with an exploratory cross-platform comparison using aggregated proprietary Facebook data retrieved through the Meta Content Library. This step should be understood primarily as a methodological experiment aimed at testing the feasibility of combining open web indexing and platform-specific aggregated data to investigate narrative circulation across online environments.

The first stage consisted of extracting narrative-specific keywords from the selected episodic narratives and using them as search queries within the Meta Content Library in order to retrieve aggregated temporal occurrence data. These outputs were then compared with Telegram observations to assess whether consistent temporal sequencing patterns could be identified across platforms.

The comparison, however, did not produce sufficiently robust or consistent results to support conclusions regarding narrative precedence, migration, or directional cross-platform diffusion. Rather than providing definitive evidence of where narratives originated or appeared first, the experiment highlighted both the analytical potential and methodological constraints of combining heterogeneous data environments for narrative-tracking purposes.

Workflow 6: CIB case study

The CIB case study represents an example of explorative, inductive research, where one observation led to another exploration. In practice, assessment of the multilingual CIB network heavily relied on persistent manual investigation, with subsequent quantitative analysis that was used to confirm or reject, at scale, what was already observed visually.

The initial observation occurred during manual data collection for the Hungarian elections monitoring. Names of political figures and parties, terms that do not translate and appear identically across languages, were used as anchor terms to identify repeated messages circulating across channels. This approach surfaced instances of the same message appearing simultaneously across multiple channels in different languages, triggering a dedicated investigation.



The initial investigation was conducted on seven directly observed channels. Each of the nine CIB indicators was assessed through the method most appropriate to it – direct visual observation for channel naming conventions and visual identity, targeted manual inspection of channel metadata for the shared contact mechanism, open-source search for external presence verification, and Telemetrio service data for subscriber growth patterns and creation dates. For the five channels whose creation dates were not available through Telemetrio, fixed-interval systematic sampling was applied, retrieving data at the 1st and 15th of each month, and inferring creation dates from the first interval returning non-zero message counts.

The network was subsequently confirmed to extend to 19 channels through a keyword-based search, independently verified by an external expert. All indicators identified at seven channels were re-assessed at 19, with no divergent results.

Quantitative confirmation of automation and synchronisation was conducted on a dataset of 222,154 messages collected across all 19 channels via the Telegram API, covering the period from 1 January to 10 April 2026. Automation was assessed through interval analysis between consecutive messages within each channel, examining whether intervals clustered around fixed values consistent with a scheduling system. Synchronisation was assessed through two complementary approaches. The first measured the proportion of channels posting within the same minute across all publishing events in the observation period. The second applied pairwise Jaccard index calculation across all 19 channels, measuring the proportion of shared posting minutes relative to total posting minutes for each channel pair. The expected Jaccard index value under conditions of independent operation was calculated at approximately 0.04, based on the observed posting frequency of approximately 12,000 messages per channel over the 100-day window.

Bibliography

- 1 Even assuming that only half of all possible minute slots are viable for publication, for example excluding overnight hours when audiences are inactive, two independently operating channels posting the same volume of 12.000 messages would be expected to post within the same minute purely by chance in, at most, 10 per cent of cases.
- 2 Smirnov's Komsomolskaya Pravda page, in Russian, can be found [here](#).
- 3 [Eto Buziashvili, Andy Carvin, Ruslan Trad & Emerson T. Brooking](#), "[How Russian propagandists spun the Trump assassination attempt](#)", DFR Lab, 18 July 2024.
- 4 Michael Colborne, "[Neo-Nazis Linked to 'Terrorist Activities' to Host Budapest Concert](#)", [Bellingcat](#), 30 January 2025.
- 5 PageRank is an algorithm that measures both the quantity and the quality of a node's connections to establish its relevance within a network.
- 6 Duncan Allen, Ognjan Denkovski & Francesca Giannaccini, "[Narratives in Motion: Measuring Cross-Platform Dissemination in a FIMI Context](#)", [Democracy Reporting International](#), 15 January 2026.
- 7 On Russian propaganda and large scale disinformation, see: Global Disinformation Unit, BBC World Service, "[The Telegram channels spreading pro-Russian propaganda in Poland](#)", BBC, 24 June 2025; and Nika Aleksejeva & Sayyara Mammadova, "[Networks of pro-Kremlin Telegram channels spread disinformation at a global scale](#)", DFRLab, 1 March 2023.
On distribution of non-consensual sexual images, see: **AI Forensics**, "[Harassment as Infrastructure: How Telegram's design enables TFGBV](#)", 23 April 2026; and Ofcom, "[Investigation into the provider of Telegram and its compliance with duties to protect users from illegal content under the Online Safety Act 2023](#)", 21 April 2026.
- 8 Kathleen Van Brempt (S&D), Parliamentary question - E-001293/2025, "[The case of Telegram and the methodology for determining VLOPs under the DSA](#)", 27 March 2025.
- 9 [Digital Services Act](#), Article 34 (1), Recitals 55 and 84.
- 10 Telegram, "[Page Verification Guidelines](#)", (last accessed 12 May 2026).
- 11 Alliance4Europe, "[Telegram takes unprecedented steps in blocking sanctioned Russian government media, but enforcement remains incomplete](#)", 13 January 2025.
- 12 European Commission, "[Communication from the Commission – Commission Guidelines for providers of Very Large Online Platforms and Very Large Online Search Engines on the mitigation of systemic risks for electoral processes pursuant to Article 35\(3\) of Regulation \(EU\) 2022/2065](#)", 26 April 2024. Recital 7: "A wide range of phenomena involving online platforms and search engines give rise to a heightened risk to election integrity. These include, but are not limited to, [...] the wider phenomenon of disinformation, the spread of violent extremist content, and content with the intent to radicalise people."
- 13 European Board for Digital Services, "[First report of the European Board for Digital Services in cooperation with the Commission pursuant to Article 35\(2\) DSA on the most prominent and recurrent systemic risks as well as mitigation measures](#)", 18 November 2025, p. 20: "Providers and CSOs have identified further systemic risks to civic discourse, including the suppression or lack of diversity of political content, polarisation, the spread of (violent) extremist content with the intent to radicalise users, the large-scale dissemination of harmful conspiracy theories, propaganda and falsely authoritative content".



Detecting Inauthentic Coordination and Foreign Pipelines on Telegram:

What We Found During
Hungary's 2026 Elections



DEMOCRACY
REPORTING
INTERNATIONAL

**DISRUPT
DISINFO**
NETWORKS