

SLOPAGANDA

The Inauthentic YouTube Network Selling Secession to Albertans

Chris Ross, Ben Steel, Zeynep Pehlivan, Mika Desblancs-Patel, and Aengus Bridgman
April 21, 2026



Key takeaways

- A network of inauthentic YouTube accounts with nearly 40 million views is targeting Albertan audiences, exploiting genuinely-held grievances and repurposing them to advance narratives that normalize the prospect of secession and U.S. annexation.
- These YouTube accounts perform an Albertan perspective, but we find no evidence to suggest any of them are Albertan. The AI avatars and paid American voice actors featured in these videos frequently mispronounce, mis-contextualize, and misunderstand the politics they cover.
- These videos contain frequent and obvious lies, drawing on real news stories to reach exaggerated conclusions designed to exploit political divisions.
- We cannot confirm this network's origin or intent, and the available evidence is inconclusive on both counts.

Points saillants

- Un réseau de comptes YouTube inauthentiques, cumulant près de 40 millions de vues, cible les audiences albertaines en s'appuyant sur des griefs réels et légitimes, qu'il récupère et détourne pour diffuser des discours qui banalisent la perspective de la sécession et de l'annexion par les États-Unis.
- Ces comptes YouTube se donnent l'apparence de représenter un point de vue albertain, mais rien ne permet de croire qu'ils sont réellement albertains. Les avatars générés par intelligence artificielle et les acteurs vocaux américains rémunérés qui apparaissent dans ces vidéos prononcent fréquemment des mots de travers, sortent les événements de leur contexte et méconnaissent la réalité politique qu'ils prétendent couvrir.
- Ces vidéos contiennent fréquemment des mensonges flagrants, s'appuyant sur de vraies informations pour en tirer des conclusions exagérées, dans le but d'exploiter les divisions politiques.
- Nous ne sommes pas en mesure de confirmer l'origine de ce réseau ni ses intentions, et les éléments disponibles ne permettent de trancher sur aucun de ces deux points.

Context and incident assessment

Alberta secession is a prominent topic in Canadian news, driven in part by a petition campaign for a referendum in 2026. Online, Alberta secession inhabits its own distinct social media ecosystem. Many figures of the movement have large X and YouTube followings, with other highly-followed, sometimes anonymous, accounts on platforms like Instagram and Facebook. Social media conversations can be difficult for the public to navigate, as democratic discourse is often interspersed with financial interests and covert influence campaigns.

We have focused on YouTube for this research because of the presence of a network of 20 inauthentic channels boosting Alberta secession and U.S.-annexation. We commend prior journalism by Colleen Hale-Hodgson and research by the Digital Forensics Lab on the topic and build on their work. This pattern is increasingly being referred to as *slopaganda* due to its low quality, highly templated approach, made easier by generative AI technology, to push political objectives through lies and exaggerations.

These channels use a mix of human actors as well as AI-generated voiceovers and avatars to deliver the scripts. Where human actors are featured, it is clear they are not the original creators of the content, as they regularly mispronounce, miscontextualize or misunderstand what they

deliver. Despite this, they often present as authentic Canadians and Albertans, using language that implies that they are integrated with the secession movement and passionate about the cause.

Because these channels offer no identifying information to real humans or organisations, nor ties to the secession movement in Alberta, we are flagging this phenomenon as a potential covert influence operation. The only person we can identify from these channels is an American voice actor based in Pennsylvania, almost certainly not the organizing force behind this phenomenon. We are concerned about the content that is deliberately inauthentic, produced by unknown actors pursuing unclear objectives.

Freedom of expression is a critical value to Canadians, one our research is committed to upholding. Advocacy for independence and sovereignty, whether from Quebec, Alberta or elsewhere, falls within that space. Our goal is to equip Canadians with the knowledge to navigate the online news ecosystem and ensure technology serves democracy. In that spirit, we believe Albertans and all Canadians deserve high quality, authentic, and fact-based news journalism. Albertan voices matter the most on topics relating to Alberta. And survey data collected by our centre earlier this year demonstrates this. Figure 1 shows the results of a question we fielded in February, 2026.

Whose views matter when discussing Alberta secession?

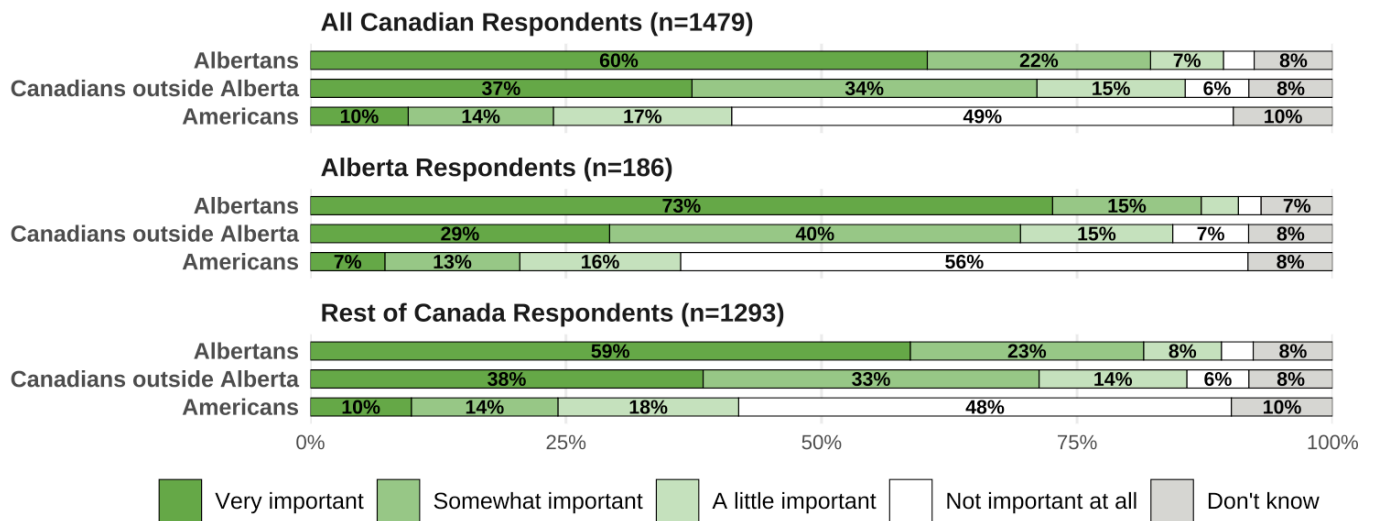


Figure 1. Whose views matter on Alberta secession?

Weighted responses to “How important is it to consider the views of the following groups when discussing Alberta secession?” MEO/Leger online panel, February 2026. Non-probability sample; see Annex for full survey methodology.

Methodology

Before analysis, we first define our scope. We began with the already-documented Canadian Reporter YouTube channel. We extracted all video titles of Canadian Reporter as a seed and searched YouTube, surfacing over 1,000 channels in results. We then measured title-level text similarity between each channel’s titles and the Canadian Reporter’s titles. By using a very high similarity threshold, we identified 35 channels that stood out for using near-identical phrasing in their video titles, suggesting coordinated or templated production rather than independent editorial decisions. From there, we manually-

reviewed the channels and excluded real, or plausibly authentic accounts and those without a dedicated focus on Albertan secession. We restrict our scope to Alberta secession-focused channels, but these selected channels are part of a broader phenomenon that deserves attention from researchers and YouTube. Finally, we restrict our analysis from May 1st, 2025 to the present as a few of these channels are repurposed accounts which feature very different content from up to 18 years ago. No channel had videos fitting the pattern of our concern before May 1st, 2025. This process resulted in the identification following 20 YouTube channels:

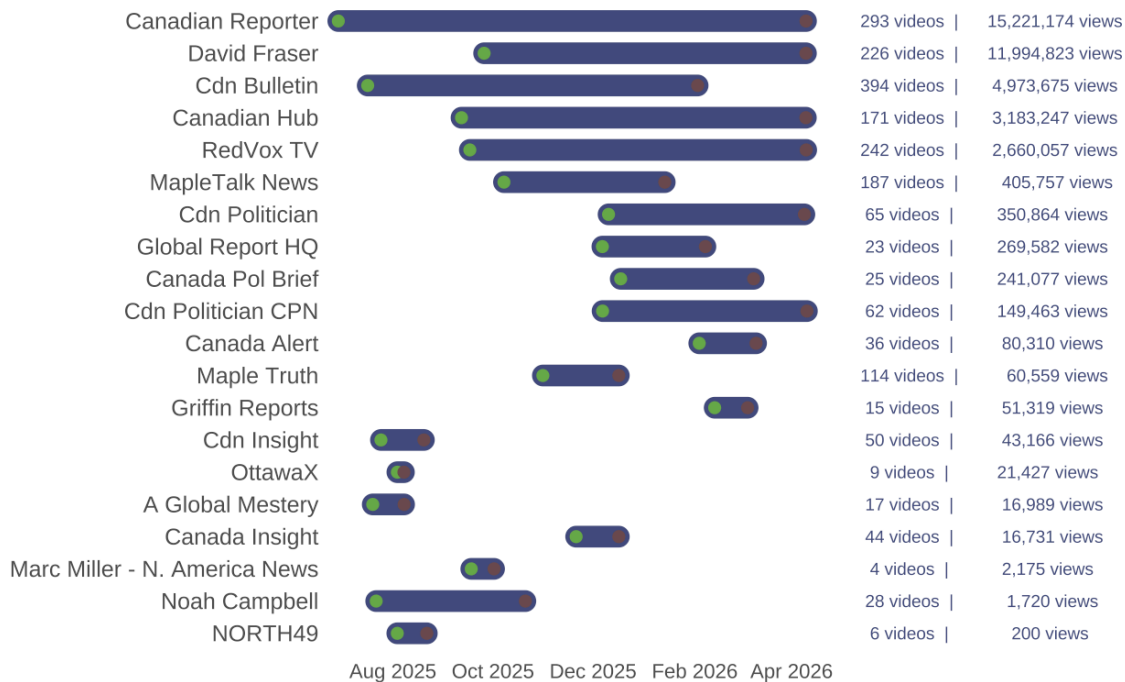


Figure 2. Inauthentic channel activity timeline

Active periods and descriptives for the 20 inauthentic channels we analyzed, with a cumulative sum of nearly 40 million views.

Research questions

What type of content do the channels cover? We downloaded and transcribed 4,474 videos across the 20 channels, segmented each into one-minute chunks (72,942 total), and used LLMs to classify each chunk for grievance framing and pro-U.S. annexation sentiment. We compare these rates against authentic Albertan secessionist channels.

Are inauthentic methods used in the creation and spread of the content? We combine qualitative review of

production red flags with computational analysis of five coordination signals (linguistic similarity, shared news clips, title overlap, templated intros, and repeated claims) alongside measures of production consistency and lexical variety.

What are the channel objectives and who is the target audience? We examine comment-level location references and situate the channels within the broader *slopaganda* phenomenon to assess plausible objectives and audiences. Origin and intent remain unresolved.

What type of content do the channels cover?

These channels use headlines-of-the-day to discuss the news from a perspective of a highly aggrieved, anti-federal government perspective. The thumbnails of the videos, likely designed to grab attention and generate clicks, give a good impression of their content and perspective. They use AI-generated deepfakes of politicians, often Premier Smith and Prime Minister Carney, with maps in the background showing some combination of western provinces as part of the US (Figure 3).

The content of the videos is often a voice over of clipped news sources, ranging from traditional media news snippets to authentic secessionist YouTube account clips. There is a high degree of similarity in clips used between channels, suggesting a template-like approach which we discuss more later. The level of work required to source, splice and edit these video segments is significant. While the content is likely disseminated across a vast amount of channels through automated processes, the original footage represents a significant production effort.



Figure 3. Most viewed inauthentic video thumbnails
Top 16 most viewed video thumbnails from the inauthentic channels.

It is worth noting that these videos include frequent and obvious lies. For example, the Canadian Reporter channel, which has over 90k subscribers and over 15 million views, opens with the following quote in their most popular video (710k views):

“The western provinces just dropped a political bombshell that is shaking Ottawa to its core, declaring moves towards the 51st through 56th statehoods with shocking authority.”

And 50 seconds later:

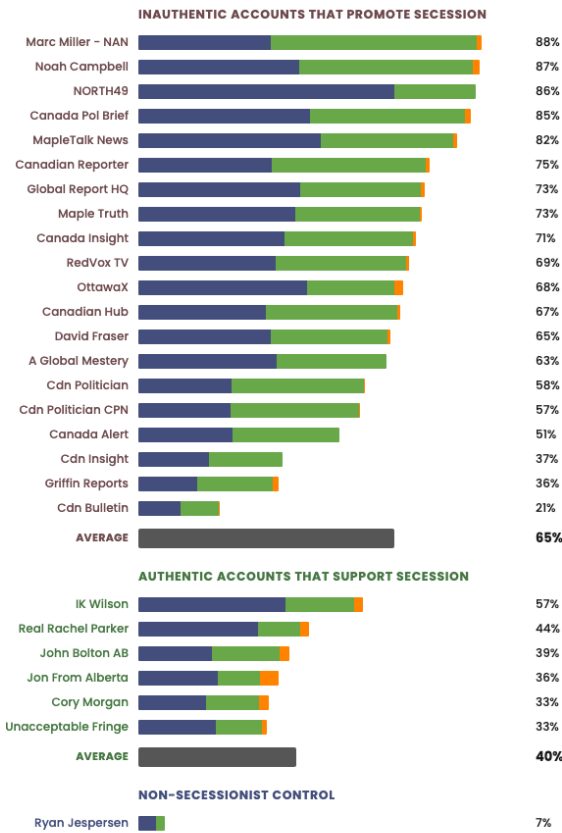
“I am witnessing decades of ignored comments, dismissed warnings, and brushed off facts finally exploding in Ottawa’s face. And here comes the punch, 65% of Albertans, 61% of Saskatchewan residents, and 58% of Manitobans openly supporting separation or even U.S. statehood. At this point, anyone who is still calling this fringe is either blind, they’re lying, or they’re desperately hoping that everyone else stays clueless.”

No published poll reports these numbers.

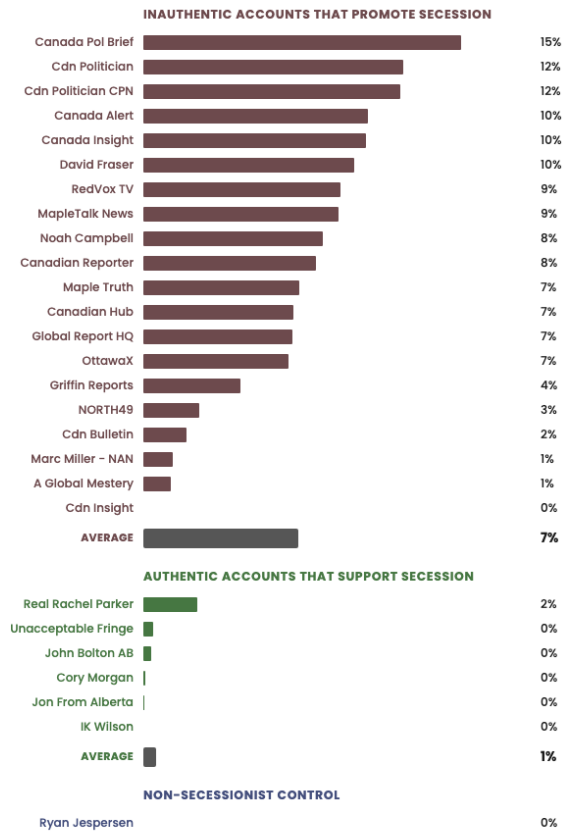
We downloaded and transcribed all the videos and used their transcripts to evaluate their content. We segmented each video into one-minute chunks then used LLMs to evaluate the presence or absence of grievance and pro-U.S. annexation framing for each chunk. Figure 4 shows these proportions for segments that include the word Alberta. We included this filter as some channels included videos on completely different topics, with only a subset dedicated towards Alberta secession content.

Our results indicate that these channels produce Alberta secession-related content characterized by substantial grievance framing alongside pro-U.S. annexation narratives. We classify grievances into three commonly recognized categories in the literature: political, economic, and cultural, consistent with leading academic research. In the inauthentic channels, grievance framing is overwhelmingly political and economic. On average, 65% of video segments in Alberta-mentioning videos contain such framing, while approximately 7% are favourable toward U.S. annexation.

Proportion of Alberta segments with grievance framing



Proportion of Alberta segments with pro-US annexation framing



Source of grievance: Political (blue), Economic (green), Cultural (orange)

Figure 4. Grievance and US-annexation framing by channel

Channel content by grievance and US-annexation favourability. Inauthentic channels are featured with dark red font, where genuine Albertan secessionist channels are featured with green font. A daily talk show, Real Talk - Ryan Jespersen, is in blue font and used as a control for a real Albertan channel that is not secessionist. Grievance: share of 1-min chunks mentioning Alberta where Alberta is portrayed as treated unfairly. Annexation (strict): share explicitly advocating joining the US / 51st state.

When comparing these findings to real Albertan secessionist YouTube accounts, we find that the inauthentic channels feature roughly 60% more grievance framing (65% of segments in Alberta-mentioning videos vs. 40% for authentic channels) and around 12 times more video segments discussing US-annexation favourably (7% of segments vs. 0.6% for authentic channels). This is with using a conservative threshold for inclusion for annexation favourable content. Below are two examples:

CanadianReporter2025: The West of Canada, Alberta, Saskatchewan, British Columbia and Manitoba have just declared their intent to form the 51st, 52nd, 53rd and 54th states of the United States if Ottawa continues blocking their independence bid. It's the kind of story that once lived in the corners of political forums, but now front and center.

RedVoxTV: Alberta, Saskatchewan, and even parts of British Columbia just dropped an earth-shattering WEXIT announcement. This isn't a drill. If we separate as an individual province or become part of the 51st state, either or, but I think we need to stand up because if we don't stand up, we're just going to be pulled over. If we did become Americans, maybe we might be better.

While authentic Albertan secessionists discuss their politics with high amounts of grievance and low favourability towards US annexation, the inauthentic channels not only amplify grievance framing but also skew the discourse toward favourable characterizations of U.S. annexation. These inauthentic accounts target Albertan audiences, co-opt their real grievances within Canada, and repurpose them to advance narratives that normalize or legitimize the prospect of U.S. annexation.

Are inauthentic methods used in the creation and spread of the content?

Having established what these channels say, we now turn to how the content is produced and spread. We combine qualitative review of production red flags with computational analysis of coordination signals across the network to assess whether inauthentic methods are used in the creation and distribution of this content.

Qualitative analysis

Before turning to computational tools, we want to give a sense of the red flags that first drew our research interest.

These channels present themselves as authentic Canadian commentary, but close listening reveals a pattern of errors that no Canadian political commentator would plausibly make. Through manual review of highly viewed videos in our channel set, we identified a series of timestamped indicators that point to the script readers not being the original creators of the content:

- Mentioning the “Atlanta prosperity project” instead of Alberta prosperity project
- Mispronouncing Regina, the capital of Saskatchewan, in a way that no Canadian would
- The voice saying the word ‘cough’ aloud instead of performing one – reading a stage direction from the script
- Calling Dallas Brodie, an Independent MLA in BC, a “he” when she is a woman. This basic error that any viewer familiar with BC politics would catch: “even the calmest politicians in BC are losing it. Independent MLA Dallas Brodie, who normally talks like *he* is reading from the hotel pillow menu...”
- Falsely claiming first-hand contact with British Columbians: “I’ve spoken to those people in Terrace, Kamloops, Fort St. John. They don’t speak the same political language as Vancouver anymore.” The script is read by a voice actor based in Pennsylvania.

These are only a small sample. Across hundreds of videos, the pattern is consistent: scripts written with surface-level familiarity with Canadian politics, performed by narrators who lack the local knowledge to catch even basic errors. The cumulative effect is a body of content that sounds like Canadian political commentary but is not produced by anyone embedded in the politics it claims to cover.

Coordinated channel introductions

The channels we analyzed used a variety of script voice-over approaches. Five use real humans that record themselves performing the content with the videos frequently cutting to news clips while they narrate videos. Others use obvious AI voice-overs and a few use AI-generated body avatars that read the scripts. We suspect the introduction of human voice actors is a response to YouTube content inauthentic guideline changes on July 15th, 2025. Figure 5 shows 4 human video narrators giving near identical channel intro videos suggesting a high degree of coordination between these channels.

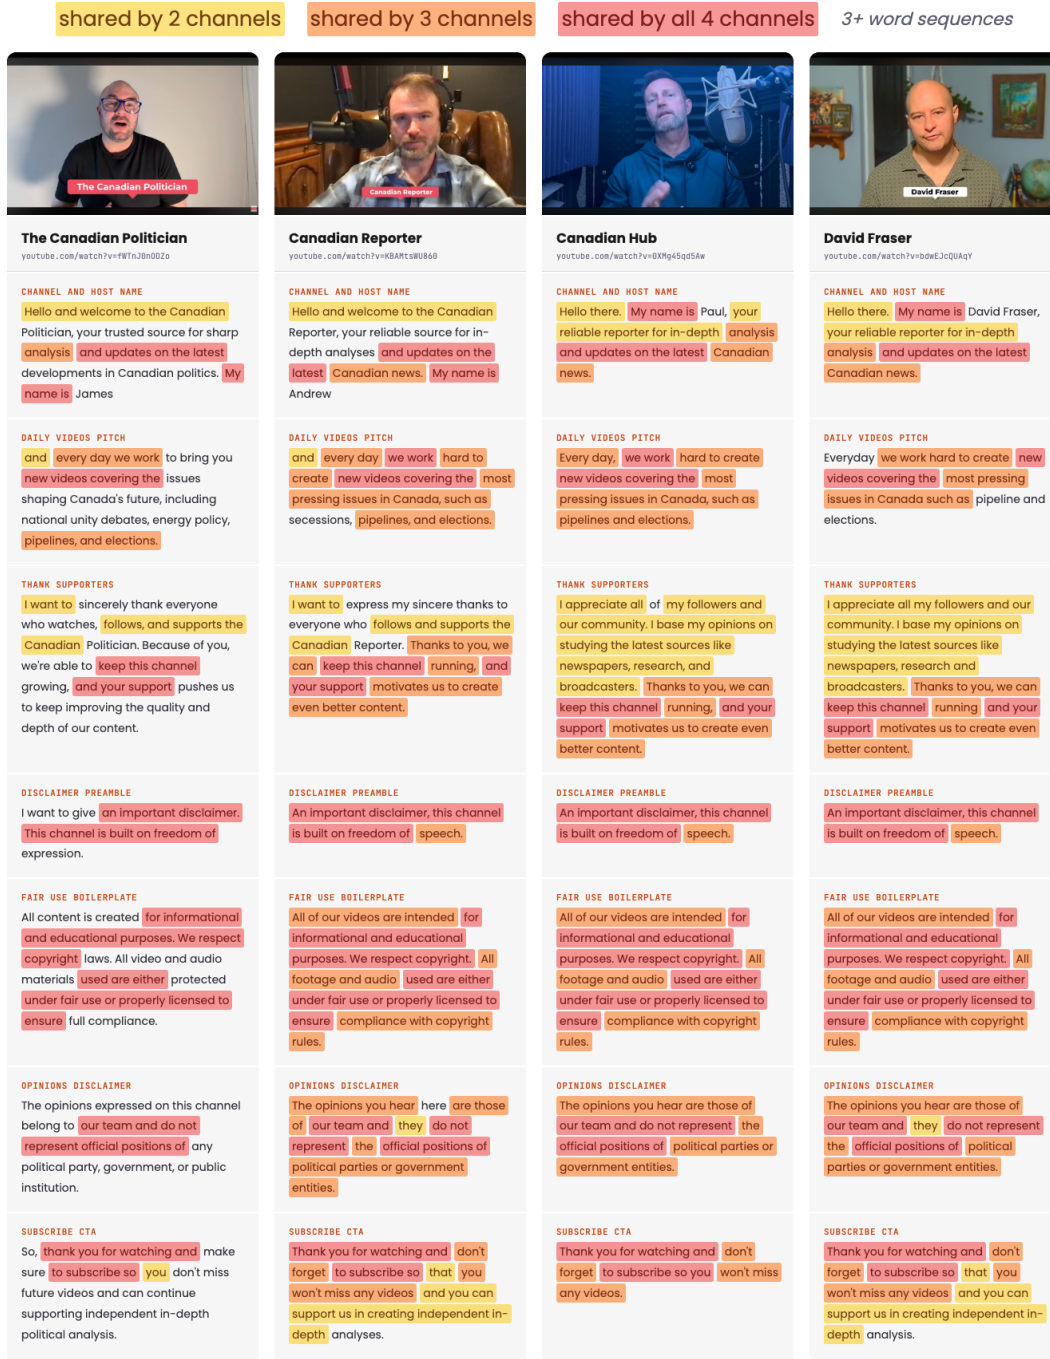


Figure 5. Coordinated channel introductions
Templated introduction videos on channels with human voice actors.

Production consistency and lexical variety

In Figure 6, we position all channels along two dimensions: vocabulary variety that measures how diverse the vocabulary is across a channel's videos and production consistency that measures how uniform each channel's videos are in length and speech rate. Channels with high production consistency produce videos that are nearly identical in duration and pacing, as though generated from the same template.

The contrast is clear. The channels we identify as inauthentic group closely together in the part of the chart associ-

ated with low vocabulary variety and highly uniform production patterns. In practice, this means their videos tend to use repetitive language and follow very similar lengths and speaking pace from upload to upload. By comparison, authentic channels are more spread out and show greater variation in both language and production style. That kind of diversity is more typical of real channels responding organically to events, opinions, and changing audiences. When combined with the other evidence in this report, it adds to the broader picture of a templated and highly systematized content operation.



Figure 6. Channel typology by vocabulary and production consistency

The 20 inauthentic channels we analyzed are featured in orange and real Albertan secessionist channels are used as a comparison, featured in blue. Bubble size proportional to total channel views.

Additional indicators of inauthentic behaviour

The production patterns documented above are not the only signal of coordination. Across these channels, we identified 11,463 accounts commenting on multiple channels (≥ 3 comments). 43% of these accounts are active on three or more channels. When we measure cross-channel audience overlap between any two channels in the network, shared commenters are overwhelmingly these multi-channel accounts, a ratio far exceeding what independent channels with shared audiences typically exhibit. The most coordinated accounts (the top 1% of this pool) appear on up to 13 channels in the network, responding to

new videos within minutes of upload and co-occurring on the same content at rates two orders of magnitude higher than activity-matched random groups.

The coordination extends to content itself. Channels share over 25,000 news clip segments that include identical audio or video reused across different uploads and repeat the same substantive claims verbatim across channels and weeks, organized around a consistent set of 20 recurring themes. Near-identical video titles are posted across multiple channels simultaneously (Figure 7). Full interactive exploration of each coordination signal is available on the project website.

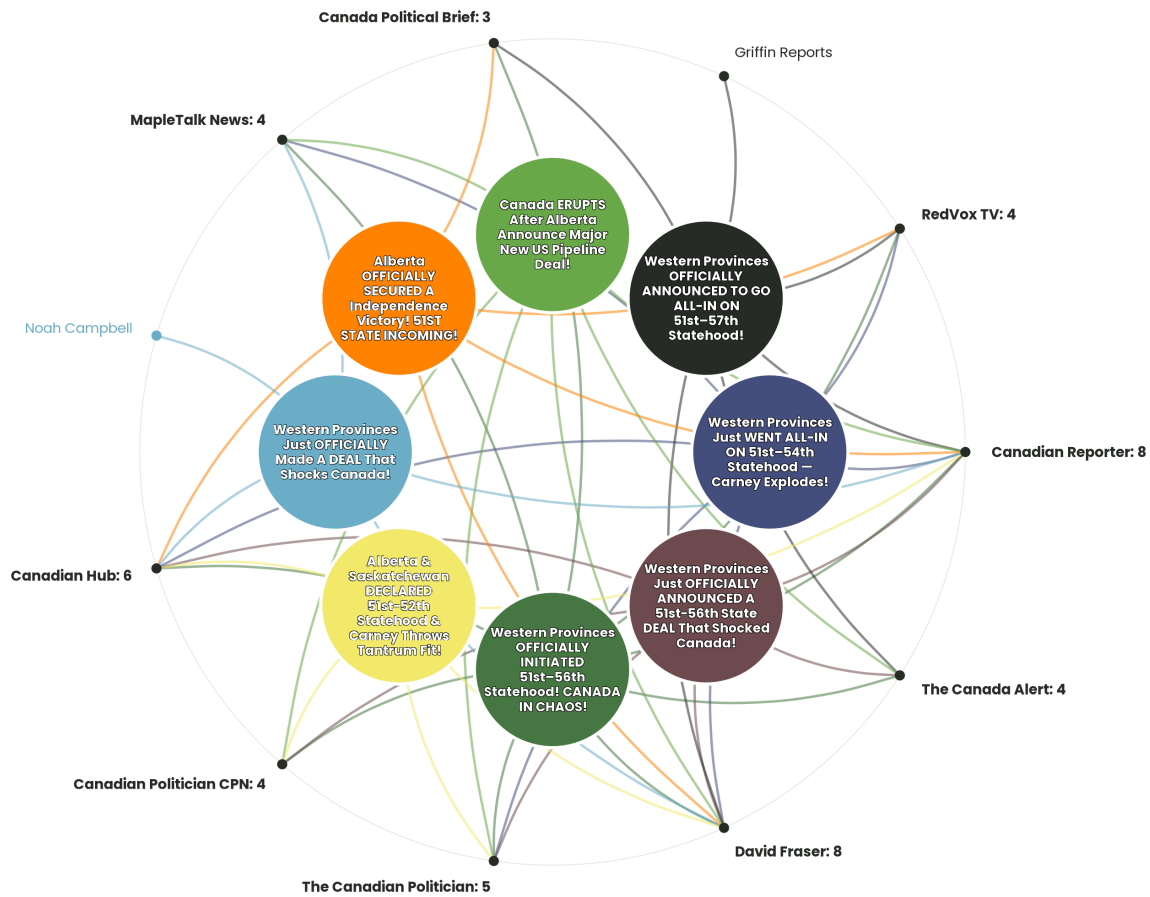


Figure 7. Shared video titles across inauthentic channels

Inner nodes are near-identical title clusters (Levenshtein similarity of at least 0.90) posted by five or more channels; outer nodes are channels; arcs link each channel to each matching title.

What are the channel objectives and who is the target audience?

Determining target audience is difficult from available data. We provide the following pieces of evidence but are not able to resolve the most pressing questions raised by this incident such as origin or intent.

A small subset of commenters (491 unique accounts, under 1% of total) voluntarily referenced a geographic location. Among these, a majority cited U.S. locations, with smaller shares referencing Alberta and the rest of Canada. We note this only as suggestive – self-reported locations in YouTube comments are not audience analytics, the sample is tiny, and such signals are easily gamed. Resolving the actual geographic reach of these channels would require platform-level data that only YouTube can provide.

These channels focusing on Alberta secession are also a part of a larger network of channels following this formulaic template of *slopaganda*. During the course of our analysis, we noticed channels scrub their videos and repurpose their account for completely different topics, such as true crime documentaries. This is a frequent practice among inauthentic content operations as platforms often determine the trustworthiness of accounts based on how long ago they were created. It can be more effective to buy an existing account and completely rebrand than to start a new one. When our team narrowed the channel scope to Alberta secession content, we passed over many other subject matters such as world politics, US politics, Iran politics, etc. Given this context, and past research on this topic, we expect the origins of these channels as broader than Alberta secession.

Alberta secession content may be featured for easy engagement returns on their content. However, the framing of the content and the strong exaggeration of grievances and US annexation content, points to political influence as a plausible objective. Co-opting authentically held grievances by real Albertans through covert methods is antithetical to transparent, democratic processes.

Concluding remarks

Our goal is to equip Canadians with the knowledge to navigate the online news ecosystem and ensure technology serves democracy. We think Albertans, and all Canadians, deserve high quality, authentic, and fact-based news journalism. This incident presents a real concern to the integrity of the Canadian media ecosystem. The channels analyzed have accumulated nearly 40 million views in the past 12 months. For context, the *Tenet Media YouTube channel*, a Russian-funded covert influence operation exposed in 2024, also accumulated 15 million views over the course of a year, though the origin of the present network remains unknown.

Alberta secession is a major political issue in 2026, with a potential referendum being held in October. We hope this research sheds light on the incident and equips civil society, governments, and platforms with the tools to respond accordingly. We call on YouTube, and social media platforms more broadly, to take seriously this abuse of the public sphere and strengthen their responses to inauthentic networks. Specifically, we call on YouTube to: (1) disclose geographic audience analytics for flagged channel networks so that researchers can determine who is actually watching this content; (2) provide transparency into whether paid promotion or ad targeting was used to direct content toward specific regions; (3) share information about the account creation and ownership history of channels identified as part of coordinated inauthentic networks; (4) extend community notes to YouTube so that viewers can flag misleading content directly (a feature that could immediately surface the fact that channels like *Canadian Reporter* are operated by paid voice actors rather than Canadian journalists); and (5) give accredited researchers API access to comment-level timestamps, account ages, and cross-channel posting patterns sufficient to distinguish organic engagement from coordinated activity. Important questions are raised by our research and we are limited in our ability to understand the origins, scope, or purposes of this inauthentic network without platform cooperation.

A note on AI in this research. AI tools were embedded at each stage of this project, from data collection and transcript classification to network analysis and prose drafting, to test what an AI-assisted research methodology could produce in a compressed timeline. Claude (Anthropic) was the primary AI tool used throughout. All code and content was reviewed, tested, and verified by the research team.

Annex: methodology

Survey

Figure 1 reports results from the Media Ecosystem Observatory (MEO) tracking survey. The survey was conducted online in February 2026 using Qualtrics with a Leger online panel. The final analytical sample consists of 1,479 Canadian adults aged 18 and older (Alberta: 186; Rest of Canada: 1,293) after excluding incomplete responses, quota overflows, screened-out respondents, attention-check failures, and duplicate panellists.

Respondents were recruited from a non-probability online panel. Because this is not a random probability sample, a margin of error cannot be calculated and results cannot be described as statistically representative of the Canadian population. Results were weighted by region, age group, and gender to match Canadian population benchmarks using iterative proportional fitting (raking), with weights trimmed to the range [0.3, 3.0].

Respondents were asked: "How important is it to consider the views of the following groups when discussing Alberta secession?" Separate items were presented for (1) Albertans, (2) Canadians outside of Alberta, and (3) Americans, each on a four-point scale: Not important at all, A little important, Somewhat important, Very important, with a Don't know option.

The survey was approved by the University of Toronto Research Ethics Board (Protocol #49009).

YouTube data

All digital trace data analyzed in this report was publicly available as of early April 2026. Transcripts of all videos are available to any journalist or researcher wishing to better understand this YouTube campaign.

Channel identification. We seeded our search with the Canadian Reporter channel, previously documented by the Canadian Press. We extracted all of its video titles and searched YouTube, surfacing over 1,000 channels. We computed pairwise title-level text similarity (Levenshtein) between each candidate channel and the seed, applying a high threshold to identify 35 channels with near-identical phrasing. Manual review excluded authentic accounts and channels without a dedicated focus on Albertan secession, yielding the final set of 20 inauthentic channels. Six authentic Albertan secessionist channels and one non-secessionist control channel were included for comparison. Our analysis window runs from May 1, 2025 to April 2026; no channel exhibited the pattern of concern before that date.

Corpus. We downloaded and transcribed 4,474 videos across the 20 inauthentic channels, producing 72,942 one-minute transcript chunks. For the six authentic comparison channels and one control channel, we applied the same pipeline.

Classification. Each one-minute chunk was classified using Qwen 3.5-27B-FP8 (run via vLLM on Alliance Canada GPU infrastructure) for (1) grievance framing – whether the segment portrays Alberta as treated unfairly, sub-classified as political, economic, or cultural – and (2) pro-U.S. annexation sentiment – whether the segment explicitly advocates Alberta joining the United States. The classification schema was designed by the research team and validated through spot-checking against manual readings of transcripts. For the content analysis in Figure 4, we restricted the denominator to segments from videos where at least one chunk mentions Alberta, ensuring the rates reflect Alberta-focused content rather than unrelated material on the same channels.

Coordination signals. Five independent signals were computed across channels: (1) linguistic similarity via TF-IDF transcript cosine distance; (2) shared news clip segments identified through audio fingerprinting; (3) title similarity via pairwise Levenshtein distance; (4) coordinated introduction videos identified through manual review; and (5) shared claims identified through HDBSCAN clustering of semantic embeddings (Harrier 270M) over approximately 300,000 extracted claims. Each signal was normalized to [0, 1] and combined into a weighted edge graph.

Production metrics. Vocabulary variety was measured as type-token ratio across each channel's transcripts. Production consistency was computed from the coefficient of variation in video duration and speech rate within each channel – channels with low variance produce videos that are nearly identical in length and pacing.

Annex: incident checklist

The following table records, against CDMRN's standard criteria, whether this case meets the threshold of an information incident warranting public reporting.

Criterion	Met	Rationale
Information Manipulation	✓	A network of 20 YouTube channels is producing news-formatted content about Alberta without disclosing authorship, funding, or location. Channels mix AI voiceovers, AI avatars, and paid actors, present as authentic Albertans, and recycle shared news-clip segments, near-identical titles, templated intros, and repeated claims across channels. Accounts scrub content and rebrand mid-study (one flipped to true-crime). These are classic markers of covert, coordinated, inauthentic behaviour.
Severity	✓	Across 4,474 videos and 72,942 transcript chunks (May 2025–April 2026), inauthentic channels frame 65% of segments in Alberta-mentioning videos as grievance (vs. 40% for authentic Albertan secessionists) and devote around 12× more segments to favourable U.S.-annexation framing (7% of segments vs. 0.6% for authentic channels). The largest channel in the set has >15M views; the network's combined reach is in the tens of millions. Content advocates the dissolution of Canadian sovereignty over Alberta.
Public Interest	✓	Alberta is running an active petition campaign toward a 2026 provincial referendum on secession. Covert amplification of grievance and annexation framing directly intersects a live democratic process of the highest political salience, and Canadians (especially Albertans) have a clear interest in knowing which voices in that debate are authentic.
Expertise	✓	MEO has the computational pipeline (transcript ingestion, LLM-based grievance and annexation coding, five independent coordination signals, network analysis) and domain expertise in Canadian political communication and platform integrity to characterize the network rigorously and at scale.

This case combines strong, multi-signal evidence of coordinated inauthentic behaviour with direct intersection of a provincial referendum on secession. The measured asymmetry between inauthentic and authentic channels (65% vs. 40% grievance framing in Alberta-mentioning videos; around 12× more pro-annexation framing) shows that the network is not merely echoing Albertan voices but systematically skewing the discourse towards grievance and U.S. annexation. Because origin and intent remain unknown, we do not attribute the campaign; we report what is observable and call on YouTube to disclose data about the network.