
CDMRN INCIDENT REPORT

FACEBOOK'S NEWS SLOP ECONOMY

Inside a foreign-run AI content-farm network monetizing fake news across Canada, the US, the UK, and Europe

Mathieu Lavigne, Mika Desblancs-Patel, David Hobson, Esli Chan, Manon Paquet, Aengus Bridgman

2026-06-22

CANADIAN
DIGITAL
MEDIA

RESEARCH
NETWORK



Executive Summary

AI has made it cheap to manufacture news. Starting from a Canadian investigation, we traced a commercially motivated network of 340 Facebook pages that floods feeds across Canada, the US, the UK, and Europe with fabricated stories to harvest ad revenue. In this incident report, we document the full international operation, focusing on its impact in Canada. We track how it reaches Canadians, what the network produces and where it misleads, and the Canadian ad dollars unwittingly funding it.

Key takeaways

- A coordinated network of 340 Facebook pages – largely operated from Vietnam – targets users in Canada, Europe, and the US with AI-fabricated stories about politics, celebrities, and sports to drive traffic to fabricated-news websites and generate ad revenue.
- Through Google's advertising systems and other automated ad exchanges, public-sector organizations and companies – including in Canada – inadvertently fund the network's business model.
- The pages command large audiences (median 11,000 followers): their posts have drawn more than 95 million reactions, comments, and shares since 2025. The network has published 355,296 posts in that period, with a peak of 3,919 posts on a single day (May 4, 2026).
- The activity is strongly coordinated: the pages are created and renamed in synchronized batches, and they recycle the same fabricated stories across the network, reusing one template for different public figures, including politicians from opposing parties.
- A fact-check of stories in the network's Canadian-news pages (5,894 posts) found that 54% were false or unsupported, typically pairing a real topic with a fabricated quote, statistic, or fictional dramatized account.

Sommaire exécutif

L'IA a rendu peu coûteuse la fabrication de nouvelles. À partir d'une enquête menée au Canada, nous avons re-tracé un réseau Facebook de 340 pages, motivé par des intérêts commerciaux, qui inonde les fils d'actualité au Canada, aux États-Unis, au Royaume-Uni et en Europe de faux récits afin de générer des revenus publicitaires. Dans ce rapport d'incident, nous documentons l'ensemble de l'opération internationale en nous concentrant sur son impact au Canada. Nous examinons comment elle rejoint les Canadiens, le type de contenu qu'elle produit, les façons dont elle les induit en erreur, ainsi que les annonceurs canadiens qui la financent à leur insu.

Points saillants

- Un réseau coordonné de 340 pages Facebook – en grande partie gérées depuis le Vietnam – cible des utilisateurs au Canada, en Europe et aux États-Unis avec de faux récits fabriqués par l'IA portant sur la politique, les célébrités et le sport, afin d'attirer du trafic vers des sites de fausses nouvelles et de générer des revenus publicitaires.
- Par l'intermédiaire des systèmes publicitaires de Google et d'autres places d'échange publicitaire automatisées, des organismes du secteur public et des entreprises – y compris au Canada – contribuent par inadvertance au financement du réseau.
- Les pages comptent de larges audiences (médiane de 11 000 abonnés) : depuis 2025, leurs publications ont suscité plus de 95 millions de réactions, commentaires et partages. Le réseau a publié 355 296 publications durant cette période, avec un pic de 3 919 publications en une seule journée (4 mai 2026).
- L'activité est fortement coordonnée : les pages sont créées et renommées par lots synchronisés et reprennent les mêmes récits d'une page à l'autre, le même gabarit étant décliné pour différentes personnalités publiques, y compris des politiciens de partis opposés.
- Une vérification des faits de 5 894 publications diffusées par les pages d'actualités canadiennes du réseau a révélé que 54 % d'entre elles étaient fausses ou non fondées. Dans la plupart des cas, elles associaient un sujet réel à une citation, une statistique ou un récit inventé.

Context and incident assessment

In the attention economy, social media platforms have become a well-known medium for profit-driven content creation, with a business model that rewards whatever captures the most attention. Users are familiar with clickbait, whether cute pet videos or outrageous rumours. The Institute for Strategic Dialogue [documented in 2022](#) how conspiracy theories have been monetized, and generative AI has since made producing such stories at scale cheaper, making content farms more lucrative.

The issue has become the subject of litigation. In February 2026, Meta [filed lawsuits](#) against deceptive advertisers in Brazil and China and against a Vietnam-based advertiser who used cloaking to run a subscription-fraud scheme. In Canada, a [class action was authorized](#) against Meta on behalf of victims of cryptocurrency scams advertised on Facebook.

As a Canadian research observatory, we began by investigating AI-fabricated content reaching Canadian Facebook users. That investigation led us outward: the pages targeting Canada turned out to be one arm of a coordinated, politically salient international network. We follow that thread in full while keeping Canada in focus, both because it is our mandate and because the clearest evidence of the network's business model – the Canadian advertising dollars funding it – is the part we can document end to end.

Amid these developments, we identify a coordinated network of 340 Facebook pages (largely operated from Vietnam, among those that disclose an administrator) that use AI-fabricated news stories about high-engagement topics in Canada, Europe, and the US to generate advertising revenue. The 340 pages we document are likely a lower bound: new pages kept surfacing throughout our investigation, as others were taken down, suggesting the network was still active and expanding as we wrote. The pages repeatedly republish the same fabricated stories while substituting their subjects (i.e., interchanging the names of politicians, sports teams, public figures, and celebrities across otherwise identical stories). These content farms generate significant engagement on Facebook, driving traffic to fabricated-news websites that are monetized through advertising. This includes legitimate campaigns from Canadian public and private organizations, meaning Canadian advertising dollars inadvertently help fund the network.

We find many similarities between this network and [documented cryptocurrency scams on Facebook](#), including

the bulk creation of pages and the use of clickbait AI-fabricated stories about high-engagement topics. The main difference is the business model, which appears centered on advertising revenue rather than financial fraud targeting citizens. Unlike a typical foreign influence operation, those behind this effort appear driven primarily by commercial rather than political incentives: across much of the network, the same fabricated stories are reused interchangeably for politicians from different parties, as a content farm chasing engagement would. A commercial motive does not make the content harmless, however. To maximize engagement the operators amplify whatever provokes the strongest reaction, and some of it – most notably the large volume of anti-immigration and anti-Muslim content aimed at UK audiences – is inflammatory. Irrespective of intent, the result is a flood of misleading, often politically charged headlines circulating on social media, particularly on Facebook.

Why this matters

For Canadians. Canadians are exposed to the network's fabricated stories directly in their feeds. In the absence of legitimate news on Facebook, these posts compete with, and can be mistaken for, credible reporting. Among users who click through, the stories are surrounded by both legitimate and questionable ads, many pressuring users to install unverified software, or pushing dubious pills or online casinos. The pages run ads targeting users aged 65 and older, [the group that is the most targeted by online scams](#).

For individuals depicted in the stories. The network routinely fabricates events about identifiable public figures – deaths, illnesses, arrests, lawsuits, and confrontations. These stories can cause reputational damage, distress families, and mislead audiences. Because templates are cheap to redeploy, the same narrative can be aimed at one person repeatedly: we observed 129 posts about the same CBC journalist falsely claiming their suspension. The risk is heightened when the operation hijacks pages that appear to have belonged to real public figures and repurposes them to publish fabricated content.

For advertisers. Advertisers who buy ad space through tools like Google AdSense may be unaware that their ads appear on pages filled with fabricated AI-generated content. This is particularly concerning for public institutions like the Government of Canada: public money ends up financing a foreign-operated network that profits from fabricated disinformation, and placing official advertising on

these sites risks lending unwarranted credibility to fabricated stories.

For society and democracy. The economy of AI-generated content poses a broader threat to the information ecosystem. Given the prevalence of these stories, users may find it increasingly difficult to distinguish them from credible journalism, eroding trust in news institutions and distorting public opinion. While the network does not appear agenda-driven, false claims about politics and politicians can mislead at scale, and we have seen many shared as if they were real reporting. Inflammatory content – such as large volumes of anti-immigration and anti-Muslim posts targeting UK audiences, and fabricated Alberta-separatism stories – could heighten social tensions.

What is this network, and how large is it?

The coordinated network of 340 Facebook pages produces content at scale that has been flooding users' feeds in Canada, the United States, and Europe. Across the network we collected 379,935 posts, with activity peaking at

3,919 posts on a single day (2026-05-04). The posts funnel users toward a large web infrastructure of 737 distinct domains that, after manual verification, we found to host almost exclusively fabricated articles, some of which have published more than 10,000 stories since the beginning of 2026. The pages have built sizable followings – a median of 11,000 followers, meaning half have more and half fewer. At the time of writing, 313 of the 340 pages were still active on Facebook, with 27 having been taken down.¹ That so many remain live is itself telling: enforcement is inconsistent rather than absent, and the pages that persist are easy to surface with the same simple searches we used. There is no ambiguity about what they are, and they should be removed.

We found that 64 % of the Facebook pages were created in or after 2025. Among pages that disclose an administrator location, 89% have administrators from Vietnam, despite biographies presenting them as based in Canada, the US, or Europe. Figure 1 summarizes the network's daily posting volume, the months its pages were created, and their administrator countries.




Figure 1: The network at a glance: recent, high-volume, and predominantly run from Vietnam. Three views of the 340-page coordinated network. (A) Posts per day ($n = 355,296$ posts from 2025-01-01 onward; the dotted line marks the daily mean of 682). (B) New pages per month, by Facebook-disclosed page-creation date. (C) Modal administrator country among the 237 pages that disclose one in Facebook's Page Transparency panel (top 10; bars are labelled with each country's share).

Taken together, these pages reach a large audience: since 2025, the network's posts have drawn more than 95 million reactions, comments, and shares (Figure 2). While most individual posts attract little attention, the network reaches this scale through large volume, with occasional stories going viral. We primarily identified these pages by

searching for content aimed at Canadian, American, and British audiences, but it has been documented that a similar network has been targeting other European countries, including France. As such, the audiences we identify here are part of a wider operation.

The identified network's reach, by content type

340 coordinated pages · 379,935 posts · 110M engagements. Each post's total engagement is counted on the day it was published; post-level country attribution.




Figure 2: The network's engagement scaled sharply through 2026 across UK, US, and Canadian audiences. Each post's total engagement (reactions + comments + shares, as recorded at collection) is counted on the day the post was published and summed across the coordinated network's posts from 2025 onward – so a day's height is the eventual reach of the posts published that day, not engagement received on that day. Each post is attributed to the audience its own content addresses: UK, US, or Canada political, or an "Other / Non-political" residual (celebrity and lifestyle filler, plus political content aimed at audiences we did not directly search for, such as Europe). Posts are classified with keyword lexicons of country-specific political figures, parties, and issues (see Appendix), prioritizing mentions of Canada so Canada-US trade content counts as Canadian even when Trump is named. Labels mark, for each strand, the publishing day whose posts drew the most total engagement.

To characterize who the network targets, we classified each post by the political audience its content addresses. Using keyword lists of country-specific political figures, parties, and issues, we labelled every post matching a country's vocabulary as political content aimed at that audience. Where a post matched more than one country, Canadian references took priority, so that Canada-US trade stories count as Canadian even when they name Donald Trump. Posts matching none of the three, such as celebrity content (the majority of the remaining content) or political content aimed at audiences we did not search for, fall into a residual "Other" category.

Across these categories, engagement has substantially increased since 2026. UK political content starts earlier and draws the most engagement, accounting for a little over half of the political content's total. US political content forms a smaller but substantial stream, while Canadian political content is the smallest of the three and scaled later. The largest share of the reach comes from the celebrity content that makes up most of what

the network publishes, accounting for nearly half of all engagement.

Why do we think these pages are part of a coordinated operation?

Our evidence that this is a coordinated network draws on two different signals: the pages frequently post near-duplicate text, and are created or renamed in same-day batches.

Near-duplicate posts across pages

Our first piece of evidence which indicates this is a coordinated network is that different pages frequently post close-to-identical stories. While the pages do not all post the same thematic content, groups of pages publish stories that follow the same template, changing only the names of the people mentioned, some adjectives, or biographical details. Identical stories were written, for example, about Nick Suzuki and Cole Caufield, and about Mark Carney and Pierre Poilievre. To identify the core coordinated cluster, we found all

page-pairs that posted at least 10 close-to-identical stories, restricted these to pages with a foreign or absent admin location, and then filtered out pages whose modus operandi we judged to differ from the main network. This procedure yields the 340-page network in-

troduced earlier, whose pages are linked by 12,611 page-pairs (pairs of pages that share close-to-identical stories). We present a visualization of this graph in Figure 3.




Figure 3: A coordinated network of 340 Facebook pages, linked by the number of near-duplicate posts and clustered using Louvain community detection. Each node is a page, sized by its total reactions. Edges link two pages that published near-identical stories, with thickness scaled to the number of near-duplicate post-pairs between them. Node colour marks their community (a consensus Louvain partition) and nodes belong to the same community if they frequently share many near-duplicate posts. Labels mark the two most-engaged-with pages in each community.

These pages separate into roughly 7 thematic clusters, with the remaining being peripheral clusters of smaller, more weakly connected pages potentially part of a wider thematic network whose posts we didn't collect. Importantly, we note that while the different page communities will generally post thematically different content, and appear to target different countries, the graph forms a single connected component – every page is reachable from every other. This means pages in different communities sometimes share near-identical stories, evidence of coordination across the countries and audiences the network targets. For example, a highly mediatized story about the disappearance of Nancy Guthrie was posted in the following four communities: "US politics & celebrity", "UK politics

& celebrity", "US politics & entertainment clickbait teasers", and "NHL hockey fandom".

Synchronized creation and renaming

Our second piece of evidence is that pages were created and renamed in synchronized, same-day batches. We identify at least 31 pages that show a consistent rebrand cycle: created under a brand name (for example, "The Carney View"), switched to an individual person's name (for example, "Gary Cannon"), and switched back to the brand name once ready for deployment. The creation itself is bulked rather than organic: groups of pages were registered on the same day, and almost all were given their first person-name rename that very same day. Most were created in late 2025 and deployed in early 2026. Figure 4

plots these rename cycles, with the same-day creation and renaming cohorts visible as vertical bands. This renam-

ing pattern plausibly aimed to avoid Facebook's account-monitoring systems.




Figure 4: Pages rebrand in synchronized waves: brand name to person name, then back. Each row is one page in the coordinated network ($n = 27$) created in 2025 or later that later returned to a name it had used before; rows are ordered by creation date. Points mark name-change events on Facebook's Page Transparency record (interim person-name flips in orange, returns to the original brand name in green); shaded bands mark dates when three or more pages renamed on the same day. A further four network pages follow the same brand-to-person-to-brand pattern but are older, repurposed accounts (created 2023-24); their multi-year histories are omitted so the synchronized 2025 cluster is not compressed. The same-day clustering is the coordination signal; the brief switch to a person's name plausibly evades Facebook's page-monitoring before deployment.

What does the network post about?

The posts heavily feature clickbait content in the typical sensationalized writing style consistent with engagement farms: all-caps headlines, emphatic and dramatic language, and real-time framings, as seen in Figure 5. Most of the posts receive very little engagement (75% garner less than 10 reactions) but a significant fraction of posts receive thousands or tens of thousands of reactions and are shared by more than 10 accounts. Figure 6 shows a map of the major topics among the coordinated posts with at least 10 reactions. Each circle represents a set of near-identical coordinated posts. Circles with the same colour indicate topically similar posts and larger circles indicate posts with a higher amount of engagement.

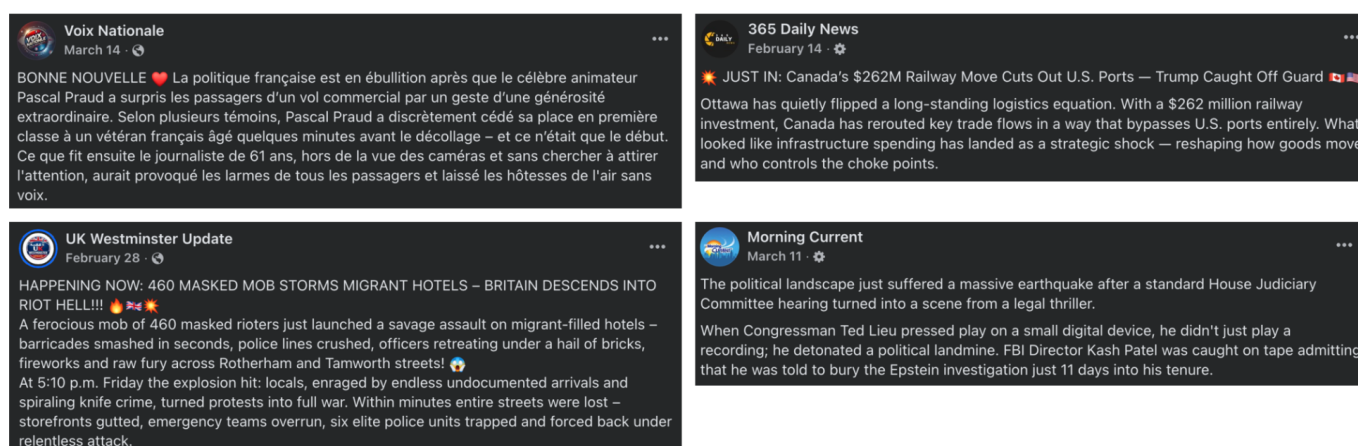


Figure 5: Sample posts from accounts targeting French, Canadian, UK, and US audiences. The post about the 460 masked mob went viral in March 2026 and received over 177,000 reactions.

By contrast, the posts targeting EU countries are mostly about France and largely focused on specific individuals, including figures such as Brigitte Macron or Pascal Praud, portraying them either positively or negatively. Besides these, most posts center around tensions within the EU, mostly with Hungary or about Ukraine. Overall, these posts represent the smallest sample in our dataset (around 5%); however, we estimate that these networks are likely much larger.

US posts chiefly revolve around fake congressional hearings—often about the Epstein files or incidents involving members of the US administration. Other prominent topics include scandals centering around TV hosts (such as Jimmy Kimmel) or other fake drama stories about other celebrities or actors. The UK posts are markedly different. They include frequent posts about scandals involv-

The major themes among the posts are diverse and are targeted towards each region's particular audience, but almost all employ shock and appeals to emotion to drive engagement. Among the Canadian content, many posts come from accounts posing as local fan pages and feature clickbait about hockey or famous celebrities. The Canadian political posts, however, largely center on themes of economic nationalism and independence – stories of Canada striking commanding trade deals, isolating the US, or Mark Carney standing up to Donald Trump. A smaller number of posts have also focused on Alberta separatism, sometimes promoting its annexation to the United States.

ing Keir Starmer but also include numerous fabricated stories of mass civil unrest: city revolts, rebellions, or military takeovers.

For the most part, these foreign accounts do not appear to contribute meaningfully to political debate or to push a coherent political agenda—the same stories are recycled across rival politicians—and the operation appears driven primarily by profit. Many pages are also found to switch between topics over time based on what is trending. For example, Dathan Thigpen's page (see the [Stolen pages](#) section below for more information about that page) first published content about US politics, switched to the Montreal Canadians when the NHL playoffs started, and began posting about the FIFA World Cup in June 2026.

As shown in Figure 6, a unique feature of the UK content is a very large cluster of posts featuring anti-immigration,

and specifically anti-Muslim, sentiment. This topic alone accounts for about 25% of all high-engagement posts. It receives both the highest average engagement (>900 likes per post) and involves the largest number of coordinating accounts (>10 accounts) among all topics. Many of these posts include fake stories of violent clashes between Muslims and Britons or Christians, or claims Muslims are attempting to impose Sharia law in Britain. Some of them,

particularly ones with the highest engagement, call for the complete deportation of Muslims from the UK as seen in the posts from Figure 7. These posts frequently mention Rupert Lowe, the leader of the Restore Britain party. While these posts often involve general themes consistent with Restore Britain's messaging, particularly around debates of immigration in the UK, they do so using completely fabricated stories.




Figure 6: Map of the major topics among coordinated posts with at least 10 reactions. Each point represents a set of near-identical coordinated posts. Each colour represents a topic and points with similar colours indicate similar topics. Larger points received more engagement.

The posts contribute to the sustained villainization of a target group that could provide the social permission for potential future violence. As Ofcom, the UK internet regulator, [concluded](#) following the 2024 riots in that country, “there was a clear connection between online activity and violent disorder seen on UK streets.” This is especially con-

cerning given the massive reach of the network. For example, one post about a “460 masked mob” attacking a hotel of migrants received over 177,000 reactions. It was eventually [debunked by Full Fact](#), but only weeks after it went live.

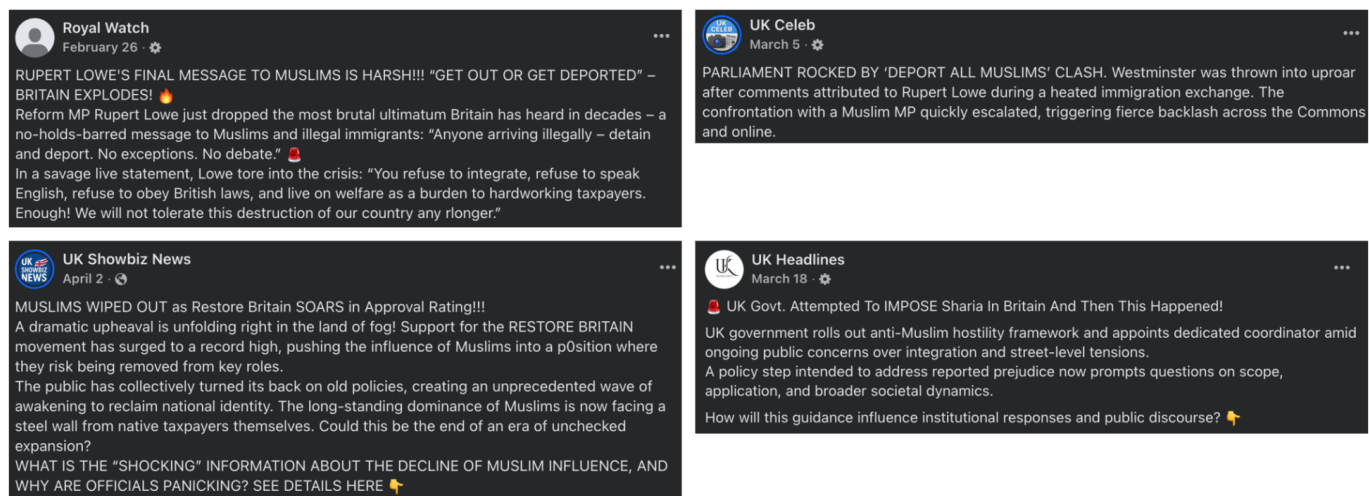


Figure 7: Sample anti-Muslim posts. All of these posts received over 6000 reactions with the two on the top row both receiving over 55,000 reactions.

How do they build audiences and drive traffic?

Because the network is funded by ads on its websites, the business model rests on attracting traffic, which requires the monetized stories to surface in users' feeds. The network uses three strategies: combining fabricated stories with high-engagement images, buying Facebook ads, and using hacked or stolen pages that already have large followings.

Fan pages and high-engagement images. Most pages present themselves as fan pages and build a following by behaving like one (i.e., posting celebrity images or sports schedules and scores, often stealing content from legitimate sources). To drive interaction, which increases visibility, the pages post images asking users to react or comment if they like or trust a public figure or organization, or to take sides in a debate using a "Like" or "Love" reaction. Some of these posts draw hundreds of reactions and comments.

Facebook ads. Many pages buy ads, usually paid in Vietnamese dong or USD, asking users to "Press here" if they trust or love a politician, public figure, or sports team. The ads generally target Canadians aged 65 and older, who receive up to 90% of impressions. Paid promotion is a cheap way to increase visibility, and even at low cost it is evidently profitable enough to be worth the investment.

Stolen pages. The network includes pages that appear to have been taken over from public figures who already had large followings. Pages that appear to have belonged to American hockey players Monique and Jocelyne Lam-

oureux and musician Dathan Thigpen – both Meta-verified – are now being used to spread fabricated stories about US politics and the Montreal Canadiens. The Lamoureux page has had more than 450 ads removed for "unacceptable business practices" between August and November 2025 (after the takeover) and was previously disabled by Meta, but has been reactivated and is now publishing high volumes of fabricated stories. Thigpen's stolen page has run ads in December 2022 targeting Mexicans with financial scams similar to those documented during the 2025 Canadian election. That accounts with this history can return to active use and keep their verified status raises questions about how consistently Meta enforces its business-practices rules.

Is there evidence of disinformation targeting Canadians?

From our investigations into page coordination, and subsequent thematic analysis of the network's content, we were able to conclude that around 35 pages are involved in posting political posts about the ongoing trade tensions between Canada and the US.

To understand their impact on the Canadian news ecosystem, we sought to characterize how truthful posts from the Canadian political news cluster actually are using *Veracity*, an automated claim-verification service developed by researchers at Mila and the Université de Montréal.² For each post the service issues a series of web searches, weighs independent sources by credibility, and returns a veracity score from 0 (no support) to 1 (well supported),

together with a written rationale. In total, the Canada-related cluster contains 6,279 unique English posts. Because the network republishes the same story with only slight variations, we collapsed the posts into near-identical stories, kept the highest-reaction variant of each, and fact-checked the resulting 1,652 distinct stories. Our assumption is that if one variant is false, its near-duplicates are too. Propagating each verdict back to every near-duplicate we get 5,894 posts.

In the table below we present these verdicts by category and show how reach and duplication compare across them. These figures describe the roughly 35-page Canadian-news cluster, not the full 340-page network. Within that cluster, we find that 55.9% of the posts are rated false or unsupported, 16.4% weak, 8.7% mixed, and 19.0% supported. Here “false or unsupported” combines stories that independent sources actively contradicted with stories for which the service could find no corroboration at all; for the dramatized, quote-based fabrications typical of this network, the latter is common. False or unsupported stories drew 2,016,503 reactions in total,

accounting for 53% of all reactions on the cluster’s fact-checked posts (3,823,546), close to their share of posts.

While the vast majority of posts are either false or weakly supported, a post’s veracity bears no meaningful relationship to either its reaction count or the number of pages that posted a variation of the story. This means that in the cluster targeting Canadian Facebook users, fake stories neither do better, nor are they spread more by the network. Pages don’t have an incentive to deliberately produce false information. Instead, the story fabrications follow consistent techniques to generate lots of plausible stories rather than random invention. Most pair a genuine theme (e.g. Canada-US trade tension, Arctic sovereignty, Mark Carney’s biography) with an invented specific: a fabricated quote, a dramatized account of a prolific figure’s reaction to an event, a fake statistic, or a flipped fact. The falseness of the story becomes a by-product of the need to create engaging content en masse drawn from real highly mediatized themes rather than a deliberate attempt at disinformation.

Table 1: Veracity of the Canadian-news community’s content. Each distinct story was fact-checked once by Veracity and the verdict applied to all its near-duplicate posts. “Reach” counts reactions on the community’s own Canada-targeted posts; copies of the same story posted by pages in other communities are excluded.

Verdict	Stories	Stories (%)	Posts	Posts/Story	Reactions/Post	Reach (%)
False / unsupported	889	53.8%	3,292	3.70	613	52.7%
Weak	273	16.5%	968	3.55	648	16.4%
Mixed	149	9.0%	513	3.44	776	10.4%
Supported	341	20.6%	1,121	3.29	697	20.4%
Total	1,652	100%	5,894	3.57	649	100%

Among the highest-reach stories that Veracity scored as false or unsupported:

- “GREENLAND “CLOSED” TO THE U.S.? CANADA-DENMARK SIGN HISTORIC ARCTIC PACT, TRUMP RESPONDS FURIOUSLY” (39,651 reactions over 14 pages)
- “3 MINUTES AGO: Trump ERUPTS as U.S. Airports FALL SILENT – 3 MILLION TOURISTS CANCEL TRIPS ALL AT ONCE” (34,100 reactions over 12 pages)
- “JAPAN PICKS CANADA – TOYOTA ABANDONS ALABAMA AS \$9B MEGAFACTORY MOVES NORTH” (33,691 reactions over 8 pages)
- “BREAKING: U.S. CORN REJECTED in a SHOCKING REVERSAL – Canada’s Premium Corn RISES to the TOP!” (29,969 reactions over 6 pages)

- “T.R.U.M.P IN TERROR as WORLD LAUNCHES MASSIVE BOYCOTT – “MAKE AMERICA GO AWAY” Goes VIRAL Overnight, Allies Turn Away...” (28,779 reactions over 2 pages)

We find that even the stories the service rated as “supported” were given scores lower than 1 because they contained a “statement appearing to be mostly true, but with some uncertainties and inconsistencies”, often due to the story containing unverifiable dramatized accounts of the headline.

In light of our recent investigation on a coordinated network of YouTube channels promoting Alberta separatism, revealed to be run by Dutch entrepreneurs, we investigated whether this network of Facebook pages also spoke about Alberta separatism. Overall, across

the entire network, we find only 298 posts talk about Alberta and only 98 about separatism. Unsurprisingly, posts about Alberta separatism are largely concentrated in the Canadian-news cluster, indicating that they are targeted towards Canadians. We don't find an overarching stance among the posts: some posts about Alberta being close to becoming a U.S. state received relatively high volumes of engagement (e.g., one post received 2,100 reactions, 769 comments, and 262 reshares), while others touting Carney's strength in standing up to separatists performed just as well – one claiming his "\$1 billion unity plan" had "destroyed" Trump's attempt to pry Alberta away drew 2,462 reactions, 331 comments, and 125 reshares. In total, posts about separatism received 32,622 reactions, and we don't see meaningful variation in reactions based on separatism stance. Running the separatism posts through Veri-fact, we find the majority are either false, unsupported, or weakly supported. Overall, the posts are simply another example of the network pushing clickbait stories with the goal of funnelling users to another website.

We also note that very few of these posts were ever fact-checked on Facebook itself. Across the coordinated network, only 88 posts carry a label from Meta's third-party fact-checking program, spread across just 43 of the 340 pages. Furthermore, in the cluster of pages posting Canada-related content, we find that not a single post was fact-checked.

What is the network's business model?

These pages generate traffic to AI-fabricated news sites, so we set out to document who advertises on those sites. We ran a ten-city advertising audit of read.usnews24s.com – a site promoted through Facebook pages like The Green Awakening – choosing one city in each Canadian province (Montreal, Toronto, Vancouver, Calgary, Ottawa, Winnipeg, Halifax, Saskatoon, St. John's, and Charlottetown) and using residential proxies so each browser appeared to be a genuine local consumer rather than a data-center bot.

In each city we set up a fresh browser and "warmed it up" the way a real local would: visiting Canadian banks, real-estate sites, local news, transit websites, and sports teams' pages, and searching expensive-keyword topics like mortgage rates, insurance quotes, and car leases. We then visited five AI-generated political clickbait articles twice per city – 100 page visits in total – recording every ad that loaded, screenshotting each, and saving the underlying creative images. We complemented this with manual visits from long-history profiles to

other fabricated-news sites, recording the Canadian-company ads shown.

Our audit confirmed that read.usnews24s.com makes its money through Google's own advertising system (Google AdSense and the DoubleClick exchange). That means legitimate Canadian advertisers inadvertently place ads on a site publishing fabricated stories about Canadian political figures, including Mark Carney and Pierre Poilievre. We captured Google-served ads from recognizable brands, including Genesis/Hyundai (a national "Signature Event" lease campaign, in Halifax and Saskatoon), Kal Tire (a "Spring Savings" promotion, in Charlottetown), La Piazzetta (a Quebec restaurant chain, in Montreal, bought through Google's premium DV360 tool), Cedar Planters Canada (Winnipeg), and Barceló Hotels (Ottawa). Manual visits to other sites in the network, such as northwavenews.com, surfaced ads from Statistics Canada, the City of Montreal, the Société de transport de Montréal (STM), Park N Fly, Expedia, Walmart, DeSerres, Metro, and the Alliance du personnel professionnel et technique de la santé et des services sociaux (APTS), among others. All of these ads were served programmatically through automated exchanges, and we have no indication that any of these advertisers knew their campaigns would appear on these sites or chose to place them there.

In sum, this foreign-operated network of disinformation sites is plugged into the legitimate Canadian advertising economy, with public-sector campaigns among the revenue sources. These advertisers may be unaware that their campaigns appear on Vietnamese-run slop sites. Across the network we also observed pop-under ads, headlines that automatically swap in the reader's city name, and monetization through providers including Google, Criteo, Taboola, MediaGo, MGID, VideoAdsTech, and AdHub Media.

Alongside these legitimate Canadian advertisers, the sites also serve questionable ads through MGID, a low-tier native-advertising network – crypto fraud, fake weight-loss supplements, and fraudulent trading platforms, geo-targeted with Canadian city names ("Trending in Toronto") and French-language ads for Quebec visitors. The same property thus runs reputable brands and predatory advertising side by side, showing the network's willingness to monetize through very different channels at once.

The network also spreads its output across multiple interchangeable sites rather than relying on any single property (Figure 8 plots this for the three sites that ex-

pose per-post dates). Publishing shifts between them over time: as one site goes dormant another surges, and dormant sites later relaunch in bulk – a pattern that

keeps the operation running even when an individual site falters.




Figure 8: Publishing migrates between the network's external sites. Articles published per day on the three external WordPress sites that expose per-post dates (the network's SPA sites carry no dates and are excluded). Counts are from a 2026-05-31 re-scrape of each site's full archive ($n = 15,243$ dated articles across the three sites); panels share an x-axis but use independent y-scales so each site's rhythm is legible. Output shifts between properties rather than running in parallel: northwavenews.com goes dormant from late March through late May while read.usnews24s.com and readnews.quantumseolabs.net surge, then northwavenews.com mass-relaunches on 22 May.

Concluding remarks

This incident provides rich evidence that some foreign actors will flood an information ecosystem with fabricated news primarily to generate revenue rather than to pursue a strategic or political objective. A commercial motive does not, however, make the content harmless. Given the network's massive reach, it floods information ecosystems with fabricated stories that can fuel disinformation narratives and contribute to social tensions. Notably, the fabricated anti-immigration and anti-Muslim stories aimed at UK audiences are particularly concerning, and could be replicated in other countries, deepening polarization.

Social media companies play a crucial role in monitoring the content they host and how their algorithms and monetization structures reward particular actors. While plat-

forms and regulators can act to reduce the visibility and monetization of this content, Canadians' ability to recognize and avoid AI-generated spam can also help make the underlying business model less profitable. Advertising systems and advertisers can also contribute to reducing incentives in this business model by redirecting their investments accordingly.

Appendix: Methodology

Page identification

Pages were identified through Facebook text-reuse search seeded from known network members and from fact-checks that flagged the network, then cross-checked against the network's external sites. We focus throughout on public Facebook Pages and exclude individual user profiles. The curated registry covers pages targeting Canada, the US, the UK, and the EU.

Close-to-identical text coordination

To identify close-to-identical story templates, we compute, for each post, the set of contiguous seven-token windows in its text (its "shingles"). The similarity between two posts is the Jaccard similarity between their shingle sets, the ratio of shared seven-token sequences to the total distinct sequences across both. Because computing this for all post pairs is expensive, we apply MinHash LSH to approximate the search for pairs above a Jaccard threshold of 0.7. Manual verification of pairs near the threshold confirmed that it captures posts sharing the same template even when names or minor details differ.

Classifying content by audience

To estimate how the network's reach divides across audiences, we classified each post by the audience its content addresses. Post text is matched (case-insensitive, whole-word) against keyword lexicons of country-specific political figures, parties, and issues for Canada, the US, and the UK; the full lexicons are listed at the end of this section. Lexicons are applied in priority order – Canada first – so that Canada-US trade content (tariffs, the "51st state") is counted as Canadian even when it also names Trump. A post matching no lexicon is treated as "Other / Non-political." The figure attributes each post individually (post-level); a page-level variant, in which a page counts as politically focused when at least a fifth of its posts are political and is then assigned to its plurality strand, gives a consistent picture but depends on that threshold, so we report the post-level view.

We validated the classification against two independent signals. First, it aligns with the data-driven coordination communities shown in Figure 3: the three political communities are 89-97% consistent with the UK, Canada, and US labels. Second, it agrees with manually recorded target country for 87% of the pages classified as political.

Because the network was assembled by searching for content aimed at Canadian, US, and UK audiences, these three strands are comparably sampled and can be compared to one another. The classification cannot capture audiences we did not directly search for (such as European or Latin American); content aimed at them falls into the "Other / Non-political" residual, and the relative sizes of the three strands should be read with this sampling frame in mind.

The full audience lexicons, matched whole-word and case-insensitive, follow.

Canada (17 terms): carney, poilievre, trudeau, freeland, jagmeet, ottawa, doug ford, danielle smith, canadian, canada, quebec, ontario, alberta, cbc, tariff, 51st state, trade war

US (47 terms): trump, biden, kamala, republican, democrat, gop, maga, congress, senate, pentagon, impeach, desantis, newsom, pelosi, hegseth, leavitt, white house, governor, vance, epstein, obama, clinton, hillary, rubio, aoc, schumer, mar-a-lago, melania, walz, fetterman, jasmine crockett, kash patel, pam bondi, tulsi, rfk, comey, marjorie taylor greene, fbi, supreme court, national guard, ice agent, ice agents, ice raid, immigration and customs, border patrol, illegal alien, illegal immigrant

UK (35 terms): starmer, keir, farage, reform uk, reform party, rupert lowe, rotherham, migrant hotel, grooming gang, westminster, downing street, nhs, bbc, labour, tory, tories, badenoch, sunak, boris johnson, tommy robinson, sadiq khan, brexit, muslim, islam, sharia, britain, british, britons, england, london, united kingdom, scotland, scottish, wales, welsh

Story fact-checking

Stories were drawn from pages that Louvain clustering identified as particularly coordinated, and whose content was predominantly about Canadian politics. We deduplicated stories by finding close-to-identical posts using the K-shingle MinHash LSH process described above. The most popular story within each template was then submitted to the *Veracity* service, which returns a veracity score, a verdict label, and a written rationale with sources. We retain the full rationale for every claim and treat the numeric score as a structured signal rather than a final verdict. Story verdicts were then propagated to each close-to-identical post.

Incident checklist

Table 2: Incident assessment

Criterion	Assessment
Inauthenticity	340 pages presenting as local; among those disclosing an administrator, 89% have administrators from Vietnam; 31 show a brand-to-person-to-brand rebrand cycle.
Targeting	Canada, US, UK, and EU audiences; ads target users aged 65+.
Coordination	340 pages sharing 12,611 page-pairs of near-identical stories, forming roughly 7 main thematic clusters with peripheral pages around them.
Scale	379,935 posts collected; peak of 3,919 posts on 2026-05-04; median 11,000 followers per page.
Veracity	Of the 1,652 distinct Canada-related stories the network’s Canadian-news pages posted, 54% rated false or unsupported and 21% supported; the rest weak or mixed.
Platform response	Few posts fact-checked or behind warning screens; verified stolen pages reactivated after removal.

A note on AI tools in this research. Claude Opus (Anthropic) was used extensively throughout this project, including drafting code, prototyping figures, and drafting prose. Veracity, an AI-based system, was also used to code the accuracy of the post content at scale. These tools make it possible to conduct a study of this depth and breadth at a speed and scale that would otherwise require a far larger team. We will continue to use them in future work and to disclose when we do.

Notes

¹We note that pages promoting fabricated content about sports teams were the most likely to have been taken down, potentially because of the greater media attention they received. See examples of media coverage: 1, 2, 3, and 4.

²Taylor Lynn Curtis, Maximilian Puelma Touzel, William Garneau, Manon Gruaz, Mike Pinder, Li Wei Wang, Sukanya Krishna, Luda Cohen, Jean-François Godbout, Reihaneh Rabbany, and Kellin Pelrine, “Veracity: An Open-Source AI Fact-Checking System,” arXiv preprint arXiv:2506.15794 (2025), <https://arxiv.org/abs/2506.15794>.