Reset.

Under the Radar:

Vast Networks of Fake Accounts Raise Questions About Meta's Compliance with the EU's New Digital Rulebook

October 2023

Executive summary

- Reset discovered an advertising network of at least *242,000* fake Facebook pages. The ads propagated by the network across the European Union promote both Russian propaganda and consumer scams, suggesting that the advertising network operates for pay.
- The network has grown exponentially since 2022, spending tens of thousands of euros on ads that violate Meta's own Terms of Service. While the Russian propaganda ads primarily targeted French and German audiences, the commercial ads promoted potentially dangerous scam products, phishing and malware to audiences in +32 countries.
- Meta has known about the network since at least September 2022, but to this day, it has failed to discontinue its malign activities, thus causing significant risks for consumers, as well as jeopardising democratic integrity in the EU. In view of next year's European elections, our findings beg the question: How does Meta intend to prevent the network from being used for targeting disinformation and Russian propaganda at millions of voters?
- Without effective mitigation from Meta, the limit on the network's size is infinite, given that new accounts can be set up by automated means at almost no cost. Meta's apparent failure to detect a basic form of automation further raises questions about the company's ability to tackle more sophisticated automation, such as swaths of content produced by generative AI.
- Reset also identified a second ecosystem of three interconnected networks engaged in similar malign activity. This ecosystem exceeds *340,000* fake pages. Failure to mitigate the risks caused by large-scale inauthentic activity on the platform raises further questions about Meta's compliance with the EU's new digital rulebook, the Digital Services Act.

Introduction

Automated and fake accounts pose critical challenges to platform integrity and security. As the biggest social media platform worldwide,[1] Facebook incubates massive numbers of automatically created accounts. Meta removes millions and sometimes even billions on a quarterly basis.[2]—according to the company's estimates[3]—often within minutes of their creation[4] with the help of an array of AI systems aimed at detecting and deactivating such accounts. Despite all efforts, the problem persists: in 2022, the platform estimated that fake accounts made up 5% of its monthly active users (MAUs), although external reports assess this number to be significantly higher, some even as high as 50% (2019).[5]

This report documents the existence of vast networks consisting of hundreds of thousands and potentially millions of Facebook pages registered as "user +" profiles. The mere scope of the networks and the common characteristics of their individual assets suggest that they were generated with the help of automated software or scripts.

Our analysis indicates these networks predominantly consist of dormant accounts, likely set up for commercial purposes. Individual pages from the networks are occasionally activated to engage in various malpractices, including pro-Kremlin advertising campaigns and scam advertising. The networks violate multiple platform Terms of Services, such as policies on coordinated inauthentic behaviour (CIB), advertising, and Community Guidelines.

Our investigation zooms in on the activities of a network of *242,000* pages launched in late 2021 and used for disseminating both Russian propaganda and scam ads throughout 2022 and 2023. This is the first-ever attempt to reveal the scale of an operation that has been ongoing for at least over a year after it was first detected by researchers. Assets belonging to the network were first discovered by the DFRLab[6] and EUDisinfoLab[7] as early as September 2022, when both organisations published reports about a pro-Russian Facebook campaign targeting EU audiences. Meta's own report[8] on coordinated inauthentic behaviour (CIB) from Russia, published in September 2022 and based on the investigation by the DFRLab, concurred with these findings, thereby confirming that the company was aware of the network's activities. According to the DFRLab's report, in September 2022 Meta had identified 703 pages belonging to the network that had been active in the campaign.

Reset's data show that by September 2022, when the reports were released, the network consisted of over 40,000 pages, almost 60 times as many as the number announced by Meta. These pages were mostly inactive, but nonetheless connected to the ecosystem. The network continued to expand after Meta's report, apparently unchecked on the platform, adding hundreds of new accounts daily throughout Q4 2022 and the first half of 2023. In June 2023, our data team estimated its total size at 242,000 pages. The network continued its advertising activities unrestricted by Meta.

1 https://datareportal.com/essential-facebook-stats.

2 https://www.statista.com/statistics/1013474/facebook-fake-account-removal-quarter.

3 https://about.fb.com/news/2021/03/how-were-tackling-misinformation-across-our-apps.

4 https://transparency.fb.com/reports/community-standards-enforcement/fake-accounts/facebook.

5 https://mashable.com/article/report-claims-half-facebook-maus-fake.
https://transparency.fb.com/reports/community-standards-enforcement/fake-accounts/facebook.

6 https://medium.com/dfrlab/russia-based-facebook-operation-targeted-europe-with-anti-ukraine-messaging-389e32324d4b.

7 https://www.disinfo.eu/wp-content/uploads/2022/09/Doppelganger-1.pdf.

8 https://about.fb.com/news/2022/09/removing-coordinated-inauthentic-behavior-from-china-and-russia.

The actor or actors behind the network are unknown. Indirect indicators such as the network's posting activity link the operation to Russian-speaking actors: 50 % of the pages have posted Russian-language content.

Our findings highlight the following:

- *Meta's approach towards the network has been selective,* ***unsystematic and retroactive,*** *limited to takedowns of individual ads or advertisers, while the ecosystem as a whole remained active and operational on the platform. The platform's algorithms failed to prevent the automatic creation of "empty and dormant assets at a high frequency",*[9] *a violative practice outlined in Meta's own policy.*
- *The company* ***continues to gain revenue*** *from political and commercial advertising launched by pages from the analysed network to this day.*
- *The pages have a few common features such as username patterns, which make them* ***stunningly easy to detect.*** *This raises questions about the effectiveness of Meta's algorithms to detect and deactivate automated accounts.*

The existence of massive coordinated networks may represent a **systemic risk** as outlined in Article 34(1)(c) of the EU's DSA: their malpractices, such as the continuous launching of foreign propaganda campaigns targeting EU audiences, may have negative effects on civic discourse and electoral processes in the EU, as well as on public security.

In the following sections, we outline the characteristics of the discovered network. We then focus on the network's most problematic activity, particularly its political campaigns and spam advertising.

In the final section, we present our findings on other massive networks seemingly operated for similar purposes and also tolerated on the platform.

We then outline Meta's general problem with detecting inauthentic accounts from coordinated networks and propose recommendations for preventing future malpractices by such actors.

I. Dissection of a coordinated Facebook network

In April 2023, Reset discovered **306 political ads** publicly available on Facebook's Ad Library that had been launched between July 2022 and April 2023 in a **coordinated campaign** to promote pro-Russian narratives targeting users from Germany and France. The ads repeated key points of the Kremlin's political agenda concerning the war in Ukraine. The peak of the campaign was between March and April 2023 when 208 political ads were launched.

A group of approximately 90 recently created anonymous Facebook pages paid for these ads. The advertisers all had unusual yet similar usernames, such as *"Delightful er3", "Great hq7", "Attractive mmg6", "Nice mg3"*. The usernames adhered to the pattern of "one positive adjective (a synonym of "beautiful) + a random string of 2–3 characters + one number". Apart from the ads they launched, the pages did not post content organically and had no followers.

To identify other pages with similar usernames, Reset's data scientists searched for accounts matching the naming convention "adjective + 2–3 letters + 1 number". This schema could accommodate thousands of username variations for every adjective by altering the subsequent letters or numbers.

9 https://transparency.fb.com/policies/community-standards/spam

We tested the detection pattern for 51 synonyms of the adjective "beautiful"[10]. Our search, completed in June 2023, revealed a network of more than *242,000* pages that had been created with those 51 words.

Over 99.5 % of the discovered network consisted of **dormant pages**. Just *750* pages (0.3 %) were **active advertisers** at the time of the analysis. The pages shared several characteristics, further distinguishing them as belonging to a common automatically generated ecosystem: similar branding identity and content strategy, patterns of posting behaviour, account type, and admin location.

1. **Common branding identity:** Most of the pages had very similar **profile photos,** typically "romantic" stock photos of Valentine's hearts, flowers, gifts, fire, or flames.

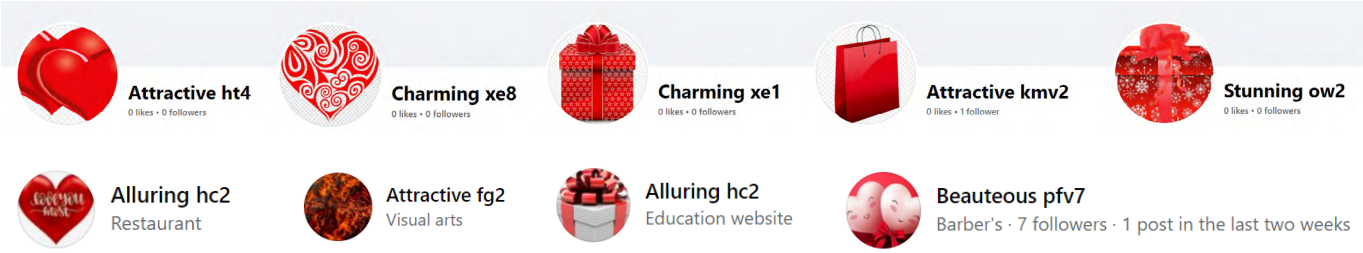

Screenshot 1: *Common visual identity: photos of hearts and gifts used as profile photos*

2. **Posting activity and content similarities.** We analysed the posting activity of the network and found that most pages had published no more than two posts. Over 99 % of the network had published between zero and three posts ***(Fig. 1)***. *70,000* pages (29 %) had never posted. The lack of posting activity indicates that those social media assets were intended for purposes other than content creation, e. g., advertising.

   The first post was typically the uploading of the page's profile photo, usually on the day of creation. Some *45,000* pages (18 % of the network) only changed their profile photo and posted no other content.

   The second post usually followed shortly after the first. For half of the network, the content strategy was eerily similar. *118,000* pages **(50 % of the analysed network)** had published a meal photo and a Russian-language cooking recipe as their second post. We found that the photos and texts of some recipe posts were taken from Russian cooking websites. The copy followed the exact same pattern, schematically enumerating the ingredients and then listing the steps for the preparation of the meal. We noticed overlaps between pages from the network that had published the same recipe

   This content consistency indicates that Russian-speaking actors are operating the network. The similarities of the posts suggest that the network is automatically generated by scripts programmed to post such content across the various assets.

10 Aesthetic, Alluring, Appealing, Artistic, Attraction*, Attractive, Beauteous, Bonny, Bright, Charming, Comely, Cute, Delicate, Delightful, Divine, Elegant, Excellent, Exquisite, Fair, Fantastic, Fine, Glorious, Good-Looking, Gorgeous, Graceful, Grand, Great, Handsome, Heavenly, Hot, Lithe, Lovely, Magnificent, Nice, Physically, Picturesque, Pleasant, Pleasing, Pretty, Property, Pulchritudinous, Radiant, Ravishing, Resplendent, Sightly, Splendid, Striking, Stunning, Superb, Winsome

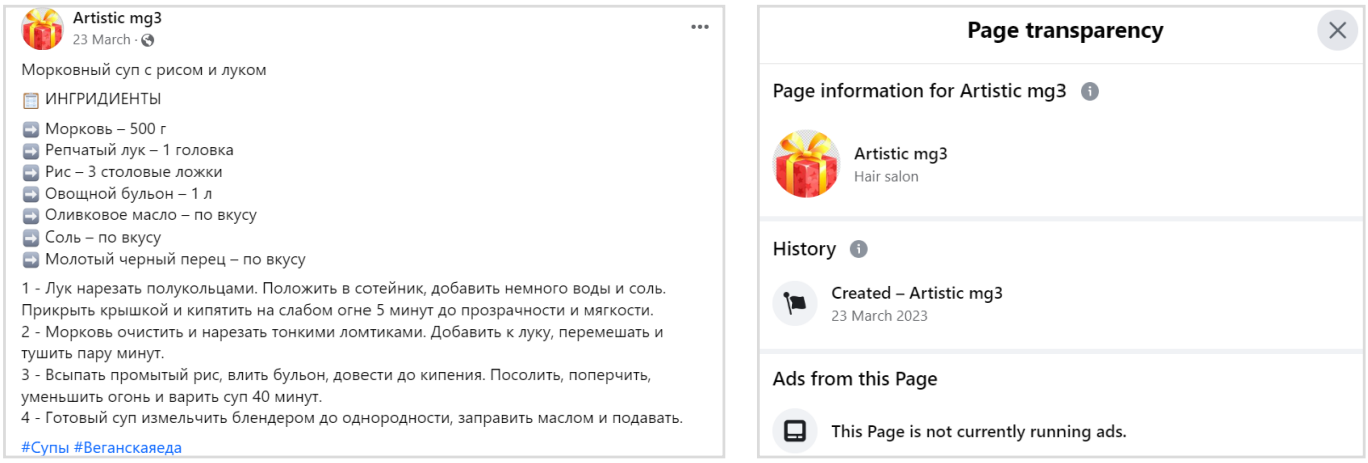

Screenshot 2: *Russian-language recipe posted by a page belonging to the network on the date of its creation*

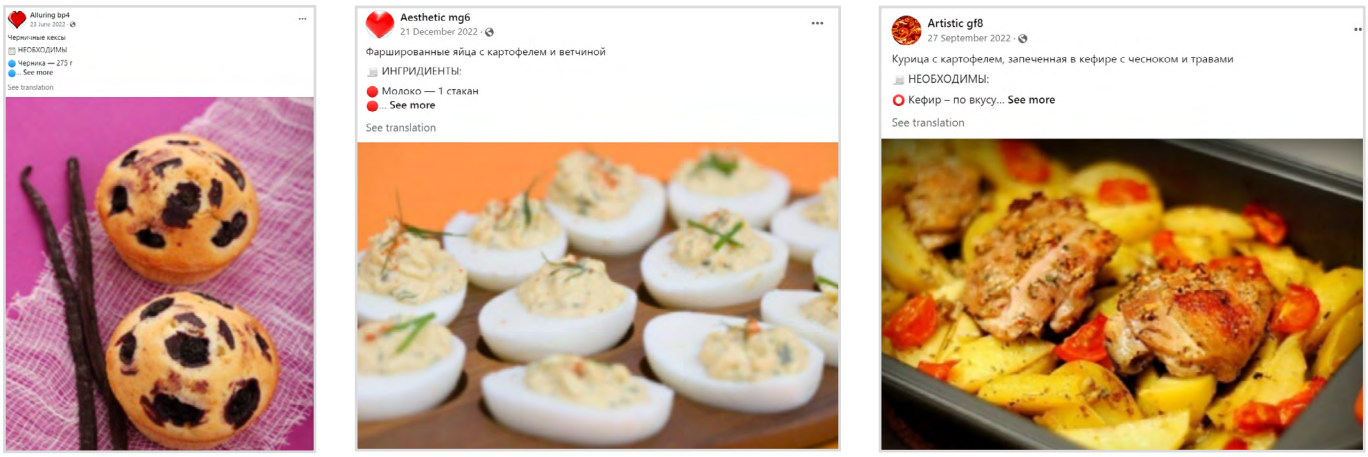

Screenshot 3: *Various Russian-language recipe posts published by pages from the network*

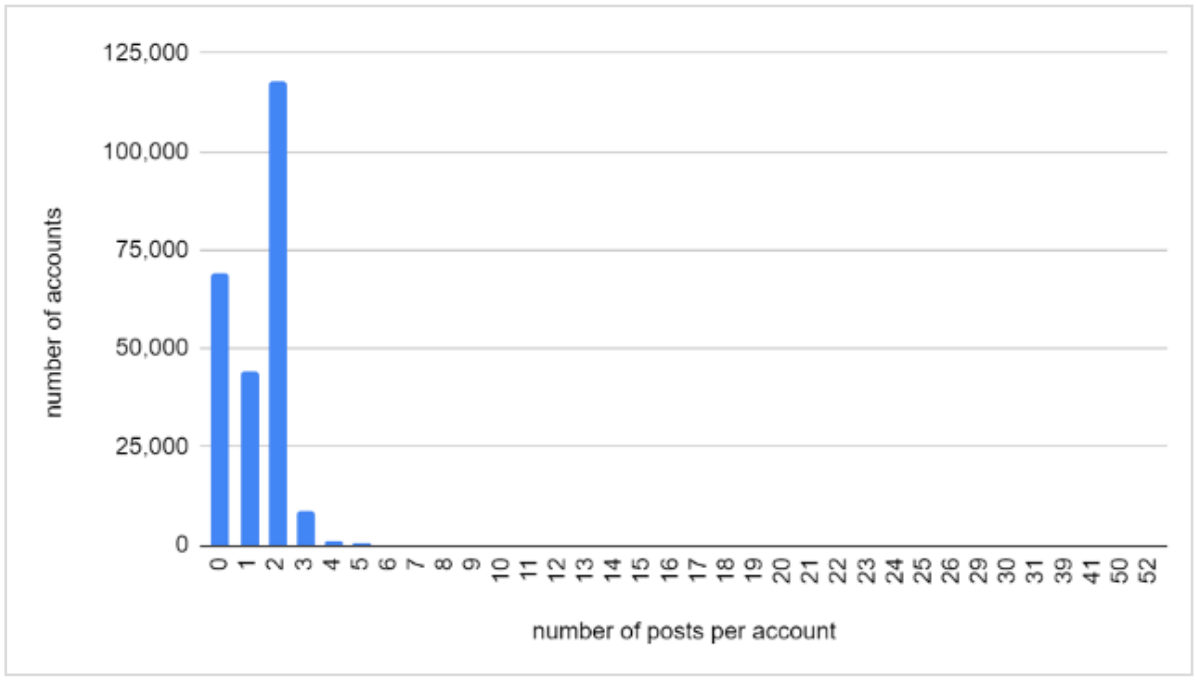

Figure 1: *Number of posts published. The most active page produced 52 organic posts in total. Most pages (~99%) published between 0 and 3 posts.*

Only a small fraction of the network continued to publish content after the Russian-language recipe post. Those posts mainly consisted of filler content, such as stock photos showing food, flowers, or landscapes.[11] 2,620 pages (less than 1 % of the network) published more than four posts. Even the most active pages posted irregularly and, eventually, would become dormant. This random posting activity is likely an attempt to keep the status of the pages active before those assets can be switched and re-used for other purposes.

3. **Creation dates.** The first pages were created in October 2021, but the vast majority of the network was launched in the months **following Russia's full-scale invasion of Ukraine.**

   The mass creation of pages intensified after June 2022 ***(Fig. 2)***, with distinct peaks around September 2022 and then later in January and March 2023. On some dates, thousands of new pages were launched.[12]

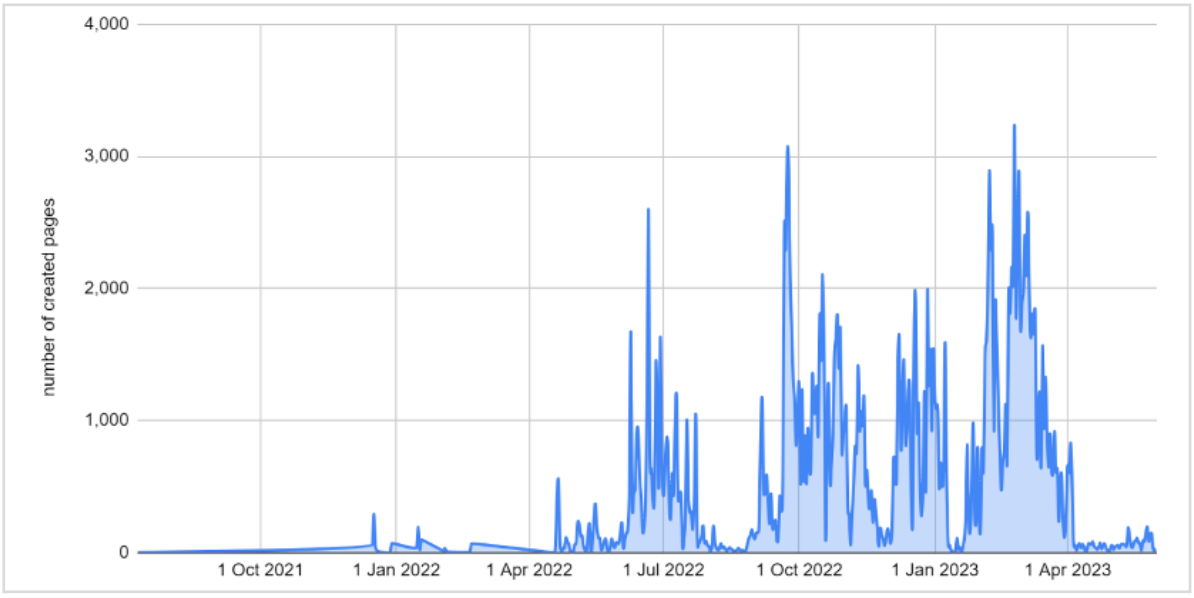

Figure 2: *Number of pages by creation date (October 2021 – May 2023)*

4. **Page followers.** We analysed data on page followers for the individual pages in the network and found that most of the pages had zero or just a few followers. The lack of followers is a direct result of their current state of dormancy. We discovered a few exceptions, i. e. pages that had accumulated 1000 or more followers despite their inactivity, which likely means that the creator had populated them with a follower network.

5. **Account types and admin location**. Almost all pages in the network were created as **Facebook "user+" accounts:** an in-between account type that is neither an individual Facebook profile nor an actual Facebook page[13] Likely because of the "user+" privacy settings, the Transparency section for most pages in the network did *not* display their admin locations. This data gap is especially problematic for the active advertisers in the network, especially the pages that paid for political ads.

11 An example: https://ghostarchive.org/archive/ct8j3.

12 An interesting coincidence is that the all-time highest peak for number of pages created on a single date fell on the 24th of February 2023, the one-year anniversary of Russia's invasion of Ukraine, when 3,200 new pages were created.

13 The fact that the pages from the network exist as "user+" profiles rather than actual pages is visible from their URLs, which contain the string "profile.php?id=", a format typically used for individual user profiles on Facebook and not for pages. Facebook pages have their own unique URLs that are different from individual profiles, and they typically follow the format "facebook.com/PageName" or "facebook.com/PageID".

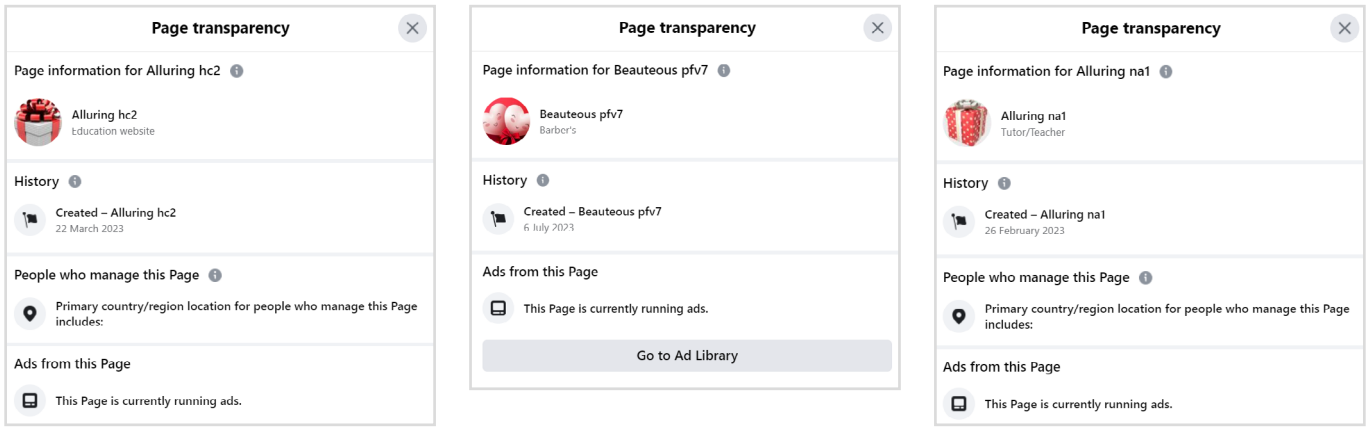

Screenshot 4: *Active advertisers with no admin location displayed in the Transparency section of the page. In some cases, the text "Primary country/ region location for people who manage this Page includes:" appears, but no country names are displayed after the column.*

Only 7 % of the pages disclosed any location. The admins of those pages were mostly based in **Ukraine** (14,300 pages), followed by **Kazakhstan** (1,300 pages). Other pages shared details of existing services and products listing Ukrainian contact details or websites belonging to actual retailers.

II. Network activity: advertising campaigns & other malpractices

As of September 2023, at least *750 pages* of the analysed network had launched ads, either political ads or commercial ads (scam ads) — potentially relevant as evidence for systemic risk under Article 34(1)(c) and (d) of the EU's Digital Services Act (DSA) respectively — with undisclosed budgets and intransparent affiliation with actual businesses.

The advertising activities continued even *after* the 25th of August 2023, when the EU's DSA regulation entered into force: we found that at least **250 scam ads** had been launched by *50 pages* from the network after that date.

Most of the political ads we discovered had been launched between January and April 2023. The earliest ad we found was from May 2022[14]; however, the total number of ads and advertisers is likely much higher, given that the network was known to investigators as early as September 2022.

Pro-Russian political campaign targeting Germany and France

In April 2023, we found that accounts from the network paid for political advertising promoting Kremlin-aligned narratives to EU audiences, mostly targeting users in Germany and France.

This campaign ran simultaneously in both countries. Third-party researchers have since corroborated our findings: Ukraine's Center for Strategic Communications exposed[15] pages from the network as participating in an advertising campaign linked to a Russian psy-op in April 2023. In June 2023, the French governmental agency VIGINUM[16] reported on the pro-Russian political campaign launched by the network in France.

14 https://ghostarchive.org/archive/avPpz.

15 https://spravdi.gov.ua/merezheyu-shyrytsya-video-de-zahid-kvapyt-ukrayinu-z-kontrnastupom-cze-chergova-vorozha-ipso/.

16 https://www.sgdsn.gouv.fr/publications/maj-19062023-rrn-une-campagne-numerique-de-manipulation-de-linformation-complexe-et.

We found **156 political ads** in **Germany** and **133 political ads** in **France** launched by pages from the network mostly between March and April 2023. In early 2023, the network also launched political ads in Ukrainian,[17] Arabic,[18] and Italian.[19]

The political ads in German and French promoted pro-Kremlin narratives in the context of the war in Ukraine: blaming NATO and the West for leading a proxy war in Ukraine, fearmongering about WWIII and the looming economic and energy crises, criticising the politics of the US and the EU, promoting fake peace initiatives in support of Russia's military agenda, etc.

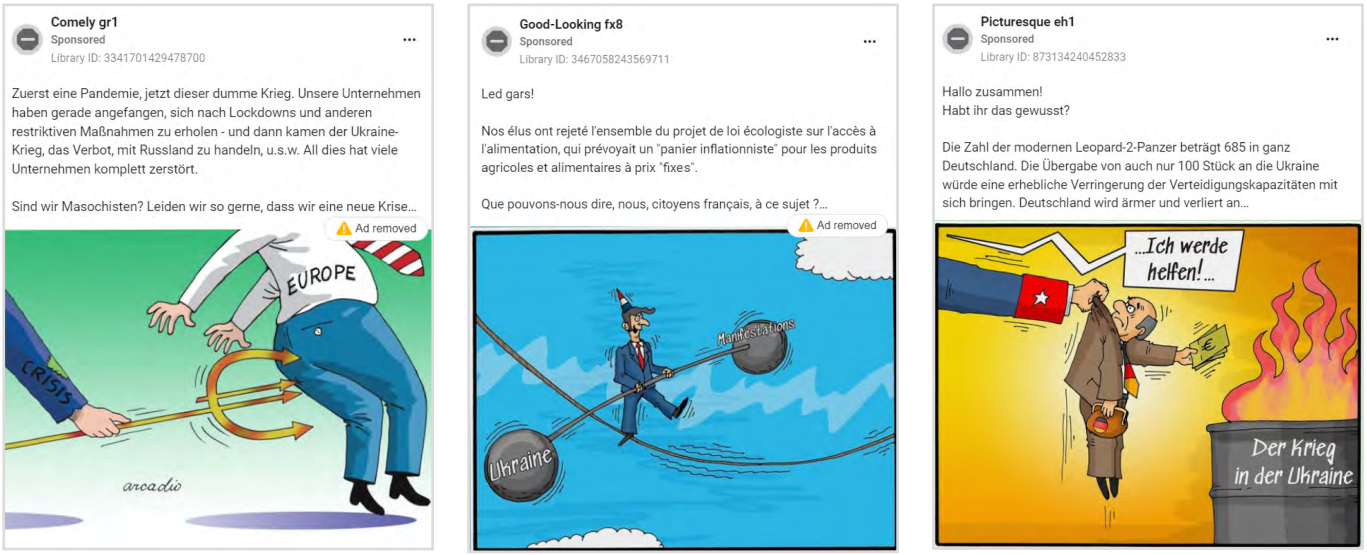

Screenshot 5: *Ads in German/French promoting narratives aligned with Russia's political agenda: fearmongering about the economic and political crises in Europe, unleashed through military help for Ukraine*

Political cartoons were the most common visual format used in these ads. Some cartoons were taken directly from Russian sources, such as the Russian Telegram channel *@VoxCartoons*. In other ads, the network used cartoons by famous illustrators, such as the Moroccan cartoonist Abdelghani Dahdouh, the French illustrator Damien Glez, etc. The network also paid for ads promoting cartoons by Chinese state media (Global Times).

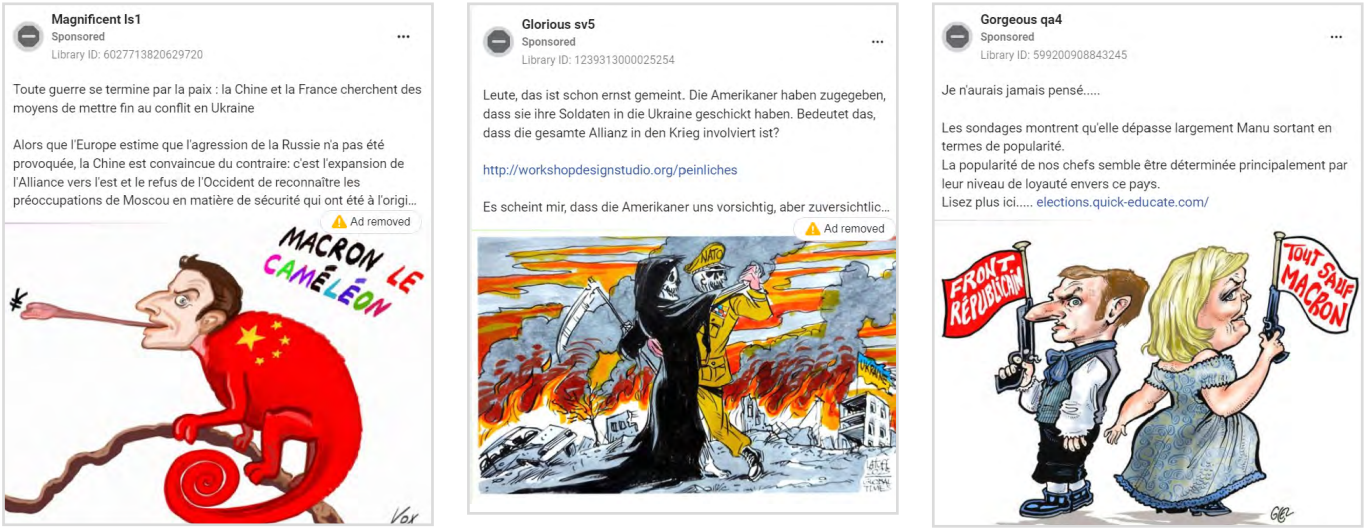

Screenshot 6: *Ads featuring cartoons by pro-Russian channel VoxCartoons (left), Chinese state media Global Times, and French illustrator Damien Glez (right)*

17 https://ghostarchive.org/archive/BBTfv.

18 https://ghostarchive.org/archive/OLGT7.

19 https://ghostarchive.org/archive/l2RL7.

We documented many instances of **cross-language promotion** targeting simultaneously users from Germany and France. The campaign included a series of **original cartoons** translated into both languages ***(Screenshot 7)*** and promoted in coordination between individual pages of the network.

Screenshot 7: *Original cartoon ads created specifically for the campaign and launched simultaneously in Germany and France: on the 31st of March 2023, the page @Exquisite xa6 first launched a cartoon ad in German referencing WWIII and blaming the West for igniting the conflict in Ukraine. Two days later, the exact same cartoon translated into French was used in an ad by another page of the network, @Property cb1.*

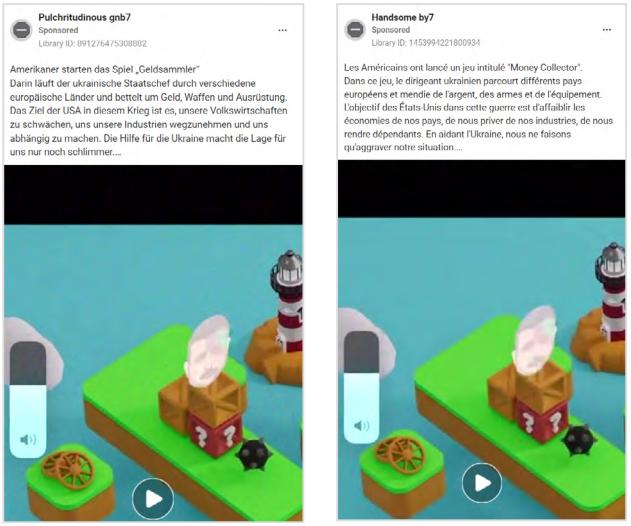

Screenshot 8*: Cross-language promotion: example of two ads with the same copy and visuals, launched simultaneously in Germany and France.*

The copy of the ads was often the same in both languages. For example, on the 9th of April, a German-language ad showed a video of a computer game ridiculing Ukrainian President Zelenskyy ***(Screenshot 8)***.[20] On the same day, two other accounts ran the same video translated into French.[21]

Some of the ads contained ***previously debunked disinformation***, already attributed to pro-Russian social media accounts. For example, on the 10th of April, three cartoon ads in German[22] parroted a story recently debunked by Snopes.com.[23] The story falsely accused Ukrainians of burning down Orthodox churches. Later on the same day, three pages from the network promoted the same cartoons in nine different French-language ads.[24]

Meta's actions against the political campaign mostly involved blocking the ads and taking down the pages of individual advertisers. Yet, some ads remained online much longer than others, and some of the pages managed to launch multiple ads before being de-platformed.

20 https://ghostarchive.org/archive/g0z8M.

21 https://ghostarchive.org/archive/LIkOU; https://ghostarchive.org/archive/Susgr.

22 https://ghostarchive.org/archive/oHsNl.

23 https://www.snopes.com/fact-check/ukrainians-burning-russian-orthodox-church.

24 https://ghostarchive.org/archive/Zeflo; https://ghostarchive.org/archive/8dUSG.

Meta profited from this campaign: the political ads cost between **1,500 and 17,689 USD** in Germany and between **1,700 and 14,611 USD** in France. These figures are imprecise because of the way ad budgets are disclosed by Facebook: many ads have a budget between 0 and 100 USD, which is displayed as <100 USD. Even with the minimum calculation of the advertising cost, Meta has received 3,200 USD from the campaign. The campaign accumulated a minimum of 1.6m ad impressions for Germany and a minimum of 876,000 ad impressions for France. In its most intense phase, the political campaign ran for nearly two months on the platform (March–April 2023), with new pages continuously joining as reinforcement to promote the content.

Commercial advertising (scam ads)

Accounts from the network also ran commercial ads containing malicious links used for phishing or information harvesting of targeted Facebook users. The links pretend to promote products or services or offer different discounts.

The scam ads used a variety of tactics to gain user traffic from Facebook, such as deceptive **redirect behaviour, typosquatting,**[25] and **misleading landing pages**. These behaviours are all prohibited by Meta's spam policy.[26]

The vast majority of the analysed scam ads referred to **newly created websites**, registered often weeks or even days before the ads were launched. We discovered 42 domains created in 2022 and 2023. Ads linking to those websites were launched shortly after that. Many of the websites were created with Tilda.cc, a website creation platform popular in Russia and Ukraine.

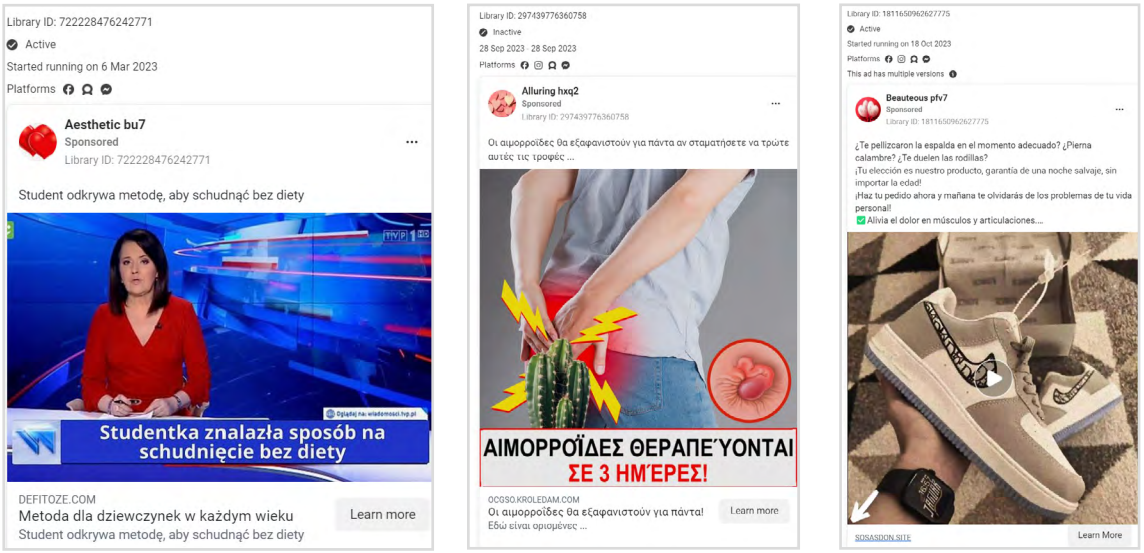

Screenshot 9: *Scam ads launched by pages from the network in English, Russian and Czech*

We found evidence of **deceptive redirect behaviour**: 63 of the ad links directed users to different websites than the websites promoted in the ads.

25 Typosquatting is a cyberattack tactic in which malicious actors register domain names similar to legitimate websites, exploiting user typing errors to redirect them to fraudulent or potentially harmful sites.

26 https://transparency.fb.com/en-gb/policies/community-standards/spam.

In 12 of the promoted links we found evidence of **domain snatching**.[27] This is a cost-effective tactic often used by scammers to obtain an online presence while retaining existing traffic from the domains. For instance, one scam ad linked to the website communicatieopleiding.com[28] a website registered in 2012 and used by a Dutch company[29] until 2014. In 2015, the company changed to a new website.[30] In 2021, the domain was acquired by a Russian actor and used to host an online store for household products.[31] In March 2023 (shortly before the scam ad was launched), the website was modified again, this time to contain job listings in German.

Our investigation identified several deficiencies of Meta's counteraction against scam advertising:

- *The platform provides no data about the cost of **commercial ads,** including scam ads. This is problematic, given that scam ads contain malicious content and proven malpractices that go against Meta's own policies.*
- *Meta takes little action against the ads and keeps no transparency record of many ads. Many of the analysed scam ads **ran until the full depletion of their budgets,** after which they disappeared from the Ad Library. Therefore, we believe the true number of scam ads launched by the network is much larger than what we could document at the time of analysis.*
- *Meta takes **no systematic action against the advertisers,** allowing some pages to continue to exist after their ads are no longer active or disappear from the Ad Library. This allows the advertisers to continue to launch new campaigns. In one extreme example, a page launched scam ads targeting 25 countries without being de-platformed.*

In total, the pages from the network have targeted at least **32 countries** with scam ads, of which 22 are EU countries.

Our investigation suggests that this is a commercial network activated for different purposes and employing various malpractices. The network exhibits a range of problematic behaviours, which Meta has neglected to action. These include accumulating fake engagement, re-branding anonymous pages to mimic businesses, and misleading audiences about company identities. These behaviours all breach Meta's Terms of Service.

- Some pages completely **re-brand themselves** as actual businesses before starting their activity with a clear intention of misleading their audiences, violating Meta's account integrity policy.[32] For example, one page[33] changed its username from *"Beauteous sc8"* to *"Natural Daily support"*, posing as a wellness online shop. Another page changed its name from *"Lovely uip3"* to *"Laptop Shop"*, before running a scam campaign about a fake laptop giveaway contest in Bulgarian.[34]
- We found evidence of **inauthentic engagement** on some pages. For example, one Russian-language recipe post received 519 likes from Facebook users based in Pakistan, India, Indonesia, and Afghanistan.[35] Another Russian-language

27 Purchasing expired domains, commonly known as "domain snatching" involves registering domain names that have not been renewed by their previous owners and have subsequently become available for registration by the general public. Those domains are usually cheaper than the newly registered ones. We found examples of websites that had previously been used by actual organisations and were later obtained by other actors to plant scam content.

28 https://www.whois.com/whois/communicatieopleiding.com.

29 https://web.archive.org/web/20140208114217/http:/www.communicatieopleiding.com.

30 https://web.archive.org/web/20160221195640/http:/communicatieopleiding.com/cgi-sys/defaultwebpage.cgi.

31 https://web.archive.org/web/20211226212223/http:/www.communicatieopleiding.com.

32 https://transparency.fb.com/policies/community-standards/account-integrity-and-authentic-identity.

33 https://ghostarchive.org/archive/xDCJ6.

34 https://ghostarchive.org/archive/egHMx.

35 https://ghostarchive.org/archive/aunRD.

recipe received 938 likes from users located in Brazil.[36] This may mean that the network likely purchases fake engagement to legitimise certain pages as authentic content creators before transitioning them to other activities, e.g., advertising.

- We found instances of **wide fluctuations in engagement**, which is a strong signal for inauthentic interaction: e.g., on one page, some posts received hundreds of likes, while others garnered no engagement at all.

III: Meta's large-scale problem with automatically created accounts

In May 2023, Reset identified **35 ads** promoting the same pro-Russian narratives with similar messages and visuals as the ads launched by the first network. However, the accounts behind these ads had usernames following the more generic naming convention *"Adjective + noun"*, with usernames such as *"Colossal page", "Rough bit", "Sneaky cat",* etc. Those usernames appeared to have been drawn randomly from lists of adjectives and nouns, likely automatically during account creation. These pages were also registered as Facebook "user+" profiles and featured a common branding identity, and some of them were also coordinated by admins located in Ukraine and Russia.

Two recent reports published by the *Ukrainian Centre for Strategic Communications* confirmed that the naming pattern *"Adjective + noun"* had been used by pages promoting pro-Kremlin disinformation before. Their reports[37] revealed a few new usernames belonging to pages from this network: *"Fresh data", "Jaded statement",* and *"Active john"*.

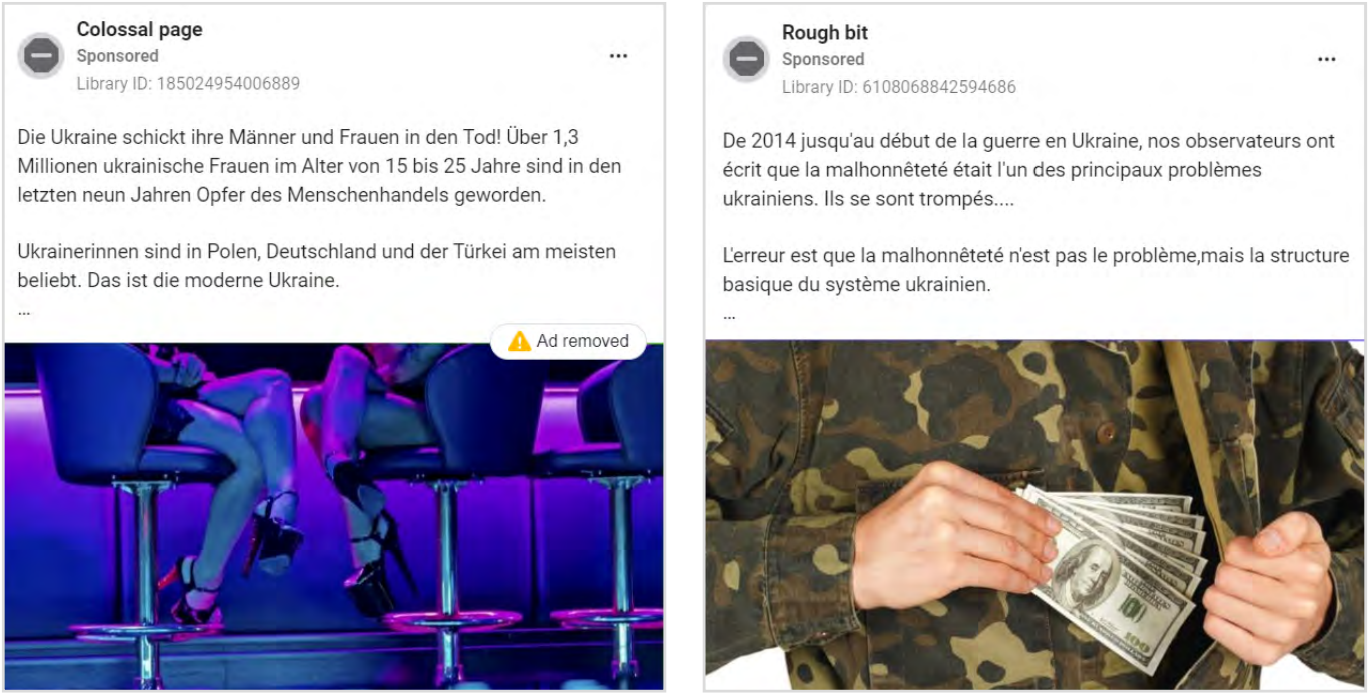

Screenshot 10: *Political ads in German and French launched by pages from another network whose usernames conform to the naming convention "Adjective + noun".*

36 https://ghostarchive.org/archive/qgLet.

37 https://spravdi.gov.ua/vorog-vypustyv-chergovu-ipso-z-novym-franczuzkym-multykom-pro-zelenskogo-ta-ukrayinu
https://spravdi.gov.ua/rosiyany-zapustyly-chergovu-ipso-z-brehneyu-pro-protystoyannya-zelenskogo-ta-zaluzhnogo.

The vast number of potential combinations arising from the simple adjective-noun pairing made it impossible to map this network comprehensively without including an unknown number of false positives.

To estimate the **potential size** of such a network, we looked at a sample of 27,500 randomly selected *"adjective + noun"* combinations. This yielded a sample of *266,500* pages matching the *"adjective + noun"* pattern and an additional sample of *54,300* pages matching a more specific naming pattern: *"adjective + noun + 5 numbers* +/- the words *"play", "video"* or *"ok"*.[38]

The search also returned accounts adhering to alternative naming patterns, such as *"adjective + noun +* the words *"DDgruopOne"* or *"DDgruopTwo"* or *"adjective + noun"* + the words *"Thanz1One"* or *"Thanz1Two"*.[39] Many of those pages displayed the name of the Brazil football player Vinicius Junior in their Intro, another indicator that they belonged to a coordinated ecosystem.

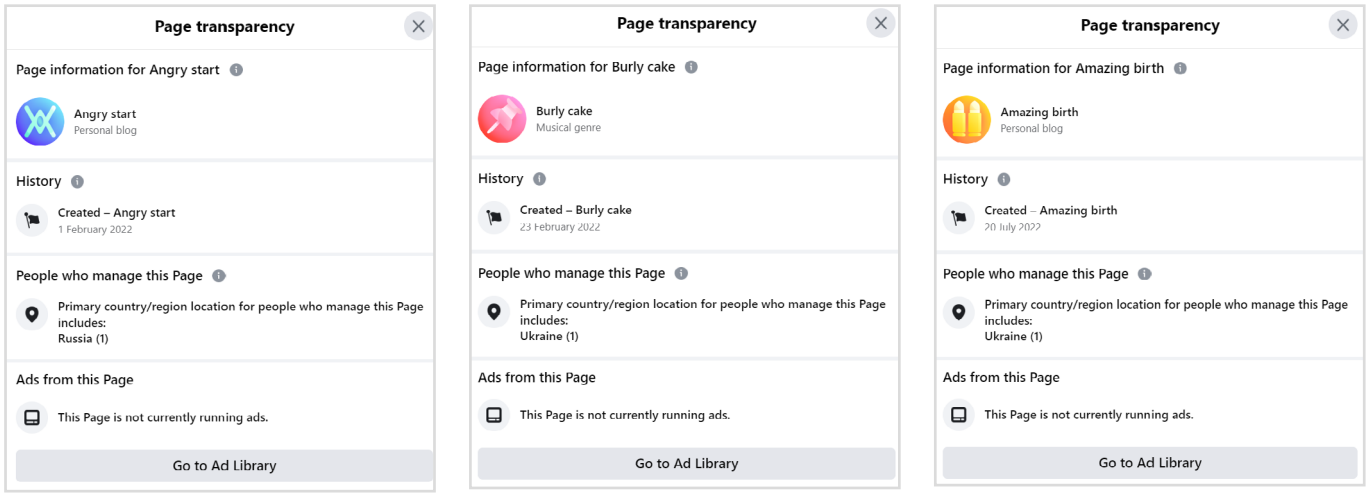

Screenshot 11: *Page usernames with the generic combination "Adjective + noun" with disclosed admin location. The pages have a similar branding identity (stock photos, gradient icons) and were created in 2022.*

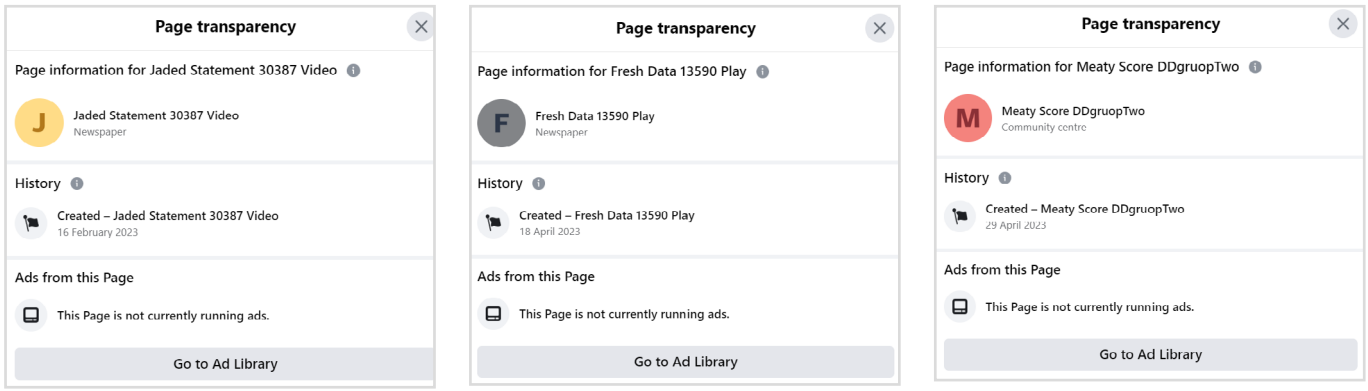

Screenshot 12: *Newly created pages with more specific usernames based on the combination "adjective + noun" with numbers and words added to the naming pattern.*

38 Examples of such pages include "Fresh Data 67141 Play", "Fresh Data 38803 Video", "Fresh Entertainmenr 28391 Video", "Jaded Statement 64481 Play", "Jaded Statement 39387 Play"

39 Examples of such pages include "Juvenile Dust DDgruopOne", "Minor Sale DDgruopOne", "Earnest Plastic DDgruopTwo", "Stylish Kitchen Thanz1One"

This mapping exercise was limited in scope, and its results are illustrative rather than definitive. They show the potential scale of Meta's problem with automatically created networks of pages. Our sample of 27,000 random adjective-noun combinations yielded over *344,000* pages.

As the naming pattern *"adjective + noun"* is common, the search likely captured genuine pages that do not pertain to the network. To measure this effect, we manually checked a sample of 400 unique page usernames adhering to the pattern *"adjective + noun"* and estimated that only 1.25 % of the sample consisted of false positives.

Conclusion and recommendations

The presence of vast networks of latent fake advertisers, unrestrained by Meta, is problematic both with regards to the company's approach to the detection of automated accounts generally and the prevention of harmful behaviour performed by anonymous pages in particular. These networks may be evidence of systemic risk under Article 34(1)(c) and (d) DSA, as well as of Meta's lack of sufficient measures to effectively mitigate these risks under Article 35 DSA.

The **creation and coordination of such networks** violate the platform's Terms of Service. Our evidence shows these networks were automatically generated, in violation of Meta's policies on account integrity and authenticity,[40] which prohibit *"the creation and use of accounts by scripted or other inauthentic means."* The mere scope of the networks is the strongest signal for automation. Numerous indicators point to coordinated behaviour — the highly specific naming patterns, the common branding identity, and the linguistic similarities in their posting activity — again in breach of Meta's policies on *"coordinated inauthentic behaviour"*.[41]

Individual pages from the networks spread political disinformation, ran scam ads, used fake identities, and exposed users to phishing and malware attacks. In response, Meta adopted a **piecemeal, reactive approach** instead of proactively detecting and de-platforming these manifestly inauthentic networks in their entirety. Removing individual deceptive ads or problematic pages does little to prevent the activation of new, empty profiles by the same actors for the same purposes.

- **Lack of ad transparency:** *Meta's efforts to curb the spread of deceptive advertising often involve not only suspending the ads but also removing them altogether from the Ad Library. In certain instances, the advertisers' profile may also be completely deleted or displayed as "Page has been unpublished or deleted." This approach hinders the accurate assessment of the impact of malicious advertising, obscures the true scale of problematic advertising behaviours on the platform, and limits the availability of evidence to researchers and investigators.*
- **Undisclosed affiliation:** *Registered as Facebook "user+" profiles, the assets from the networks enjoy privacy settings on a level of anonymity inappropriate for accounts engaging in malicious practices — e.g., the Page Transparency section discloses no information about admin location for the majority of the pages.*
- **Undisclosed revenue:** *The exact cost of many **political ads** launched by the analysed network is displayed as just a rough estimation (e.g., < 100 USD). This makes it impossible for researchers to collect precise information on how much Meta has profited from the advertising. Meta also does not disclose the budgets of **commercial ads,** including scam ads exposing users to various risks (phishing, identity theft, malware, etc.), which is a serious transparency issue given their problematic content.*

40 https://transparency.fb.com/policies/community-standards/account-integrity-and-authentic-identity.

41 https://transparency.fb.com/policies/community-standards/inauthentic-behavior.

The actors behind networks of such scope typically implement various **techniques to avoid detection**, including rotating IP addresses, running on multiple machines or using distributed services, and introducing delays between account creation. However, by applying more robust algorithmic detection techniques, Meta should be able to identify these large clusters of nearly identical accounts. This especially applies to cases where individual assets have already been exposed.

- *One recommendation to the platform would be to introduce more robust pattern recognition algorithms, which would allow the detection of pages with similar usernames and behavioural signatures.*
- *Strengthening the platform's content analysis capabilities by machine learning models that identify repetitive or suspicious content (such as the same cooking recipe being used as the first post on various pages from the studied network) is another step that would allow the detection of clusters of accounts.*
- *Another recommendation is to apply more robust verification measures for accounts that run ads. This may include stricter validation of phone numbers, email addresses, or other forms of identification to minimise the chance that automated accounts are used for advertising.*

Given the potential negative effects of such networks on civic discourse, democratic processes and even public security, our investigation underscores the need for consistent and continuous action against such actors. This is particularly poignant now, given that 2024, with the EU election and national elections in over 45 countries, will be a pivotal time for democracy.