Monitoring the Implementation of the Code of Conduct on Disinformation through Structural Indicators: Gaps, Evidence, and Policy Implications

Teresa Weikmann, Guy De Pauw, Marina Tulin, Michael Hameleers, Claes de Vreese

Executive Summary

Under the **Code of Conduct on Disinformation**, several social media platforms have committed to mitigating the spread of mis- and disinformation on their services. In practice, however, a gap remains between the demand for information on what platforms actually do – namely, the structural indicators of their actions – and what platforms disclose in their self-reports. In this white paper, we address the challenge of monitoring the implementation of the Code of Conduct across social media platforms through three complementary approaches:

1. An **assessment** of platform signatories' **self-reports**
2. A review of existing **political communication literature** on the prevalence of disinformation on platforms, and
3. A **case study** examining the prevalence of disinformation and common **narratives in Dutch-language regions**.

Overall, we find that platforms disclose only limited information in their self-reports. Our own analysis is consistent with findings from the academic literature and shows that while the overall prevalence of disinformation is relatively low, it nonetheless persists and reveals recurring problematic narrative structures. Based on these findings, we recommend incorporating the dimension of **false and misleading narratives into the set of structural indicators**. Doing so could incentivize platforms to provide more detailed and substantive self-reporting on the content they identify as problematic and the measures they take in response. Furthermore, such an indicator would be relevant for media professionals, including fact-checkers, to understand which false narratives are circulating online and how best to counter their spread and potentially harmful effects.

Co-funded by the European Union

Introduction

Social media platforms are at the center of debates about the spread of mis- and disinformation, that is, the inadvertent and intentional dissemination of falsehoods, respectively. They are considered a potential breeding ground for such content, given their unprecedented control over information visibility and amplification, coupled with their massive scale and opaque algorithms. To this day, researchers are grappling with the question of how much disinformation is actually present on platforms, where it comes from, who is exposed to it - and what platforms are doing to mitigate its spread. Yet, these are exactly the kind of metrics – or "structural indicators" – that platforms have committed to self-reporting under the European Union's **Code of Conduct on Disinformation** (previously the *Code of Practice*), a co-regulatory framework aimed at increasing transparency and accountability. This means that although platforms should report the metrics necessary to assess mis- and disinformation on their services, a substantial gap persists between the demand for such data and what is disclosed.

In this paper, we first monitor the implementation of the Code of Conduct on Disinformation across EU-countries by analyzing the most recent platform self-reports under the code (covering the period 1 January – 30 June 2025) submitted by all active social media platform signatories (LinkedIn, YouTube, TikTok, Meta), highlighting the limited scope and depth of the information disclosed. Second, we synthesize recent empirical communication science research on how widespread mis- and disinformation is across major platforms, thus assessing platform compliance with the Code from an independent researcher's perspective. Third, we present the Netherlands and Belgium as a case study, illustrating the prevalence of disinformation *narratives* in the Dutch online environment during the same time frame as covered in the self-reports, thereby zooming in on the first structural indicator (prevalence of disinformation).

Overall, this paper contributes to a deeper understanding of how the Code of Conduct on Disinformation is implemented across platforms, identifies key challenges in its execution, and outlines areas for improvement moving forward. Most crucially, we assess the potential gap between the informational demands of structural indicators and the actual supply of valid and reliable data on disinformation narratives across digital information settings. The adequate monitoring of the Code requires access to valid and reliable data on disinformation narratives, which currently does not result from the reporting of social media platforms. We therefore explore how independent (academic and policy) researchers may supplement data and offer data that can be used to monitor the performance of platforms. Moreover, we explore the feasibility of such approaches, and herewith the sustainability of monitoring and auditing tasks by independent researchers vis-a-vis platforms themselves.

Co-funded by the European Union

Structural indicators under the Code: A look into platform self-reports

In this section, we use the term disinformation (rather than misinformation), which is preferred by policymakers, given that emphasizing intent and harm provides a stronger basis for policy intervention (Bleyer-Simon & Reviglio, 2024). In line with this, the Code of Conduct on Disinformation (CoC) is designed to function as a co-regulatory framework for – amongst others – very large online platforms (VLOPs) to limit and safeguard against the spread and impact of disinformation (European Commission, 2025). As such, it is integrated into the framework of the Digital Services Act (art. 45 DSA), making it a relevant benchmark of compliance. Most major platforms have reduced or renegotiated their original commitments in key areas while transitioning from a Code of Practice to a Code of Conduct, often presenting a more ambivalent stance (see also Böswald, 2025). Currently, four social media platforms (or VLOPs) are signatories: LinkedIn, Meta (including Facebook and Instagram), TikTok and YouTube (as part of Google). These platforms have – to a certain extent – committed to taking relevant measures to prevent the spread of disinformation on their services. In addition, under Commitment 41, they opted to assess the effectiveness of their measures by reporting on so-called structural indicators, which reads as follows:

> "Signatories commit to work within the Task-force towards developing Structural Indicators, and publish a first set of them within 9 months from the signature of this Code, and to publish an initial measurement alongside their first full report. To achieve this goal, Signatories commit to support their implementation, including the testing and adapting of the initial set of Structural Indicators agreed in this Code. This, in order to (sic) assess the effectiveness of the Code in reducing the spread of online disinformation for each of the relevant Signatories, and for the entire online ecosystem in the EU and at Member State level. Signatories will collaborate with relevant actors in that regard, including ERGA and EDMO."
> (European Commission, 2025, p.47)

Figure 1. *Commitment 41 under the Code of Conduct on Disinformation, referring to structural indicators.*

To make these structural indicators (SIs) measurable, a group of experts developed a set of relevant metrics. These encompass:

- The prevalence of disinformation (SI1),
- its sources and origins (SI2),
- audience thereof (SI3),
- its demonetization (SI4), and
- platform investments in fact-checking (SI5) (see Nenadic et al., 2024).

Co-funded by the European Union

The official webpage of the Code, also called the Transparency Centre, hosts biannual self-reports in which platforms describe the measures and efforts undertaken to comply with their various commitments. However, these transparency reports largely focus on so-called **service-level indicators** – qualitative measures that capture platform's actions to support fact-checkers, users, and researchers (Mündges & Park, 2024). To examine what kind of information is provided regarding **structural indicators** specifically, we conducted a qualitative analysis of all most recent VLOP reports, published in September 2025 (covering the period 1 January–30 June 2025). Our analysis proceeded in two steps. First, we reviewed the information that signatories provided under Commitment 41, which directly addresses structural indicators (see above). Second, we searched the reports for any quantitative data offering insight into these indicators, such as tables or figures that displayed amounts or engagement with misinformation or any metrics regarding its demonetization. Importantly, the reports usually use the term *misinformation*, given the difficulty of proving intent that is typically covert and not identifiable through hard content indicators that can be accessed without interpretation or contextual evaluations. Hence, even though the Code addresses *disinformation* in principle, we use the term *misinformation* when referring to information taken directly from the self-reports. Overall, this analysis aimed to answer the following research question:

Method: We conducted a qualitative analysis of the most recent VLOP transparency reports released in September 2025 in two steps: First, we examined all disclosures submitted under **Commitment 41**, which pertains specifically to structural indicators. Second, we screened the reports for any **quantitative information** relevant to these indicators, including tables, figures, or metrics that referenced misinformation volumes, prevalence, or impacts.

RQ1: *To what extent do platform signatories report on structural indicators in their most recent biannual transparency reports (covering 1 January–30 June 2025)?*

Results

LinkedIn

The LinkedIn report is 67 pages long, making it the shortest of the four VLOP reports. At the beginning of the report, under Commitment 2, LinkedIn states its policy: "Do not share content that is false, misleading, or intended to deceive" (LinkedIn, 2025, p.8). The platform provides no information under Commitment 41 which regards developing, implementing and reporting on structural indicators, indicating that it has "not subscribed" (p.51). Notably, LinkedIn (as part of Microsoft) did not subscribe to the fact-checking commitment when transitioning from the voluntary Code of Practice to the Code of Conduct. Regarding quantitative data, LinkedIn reports the number of ads

Co-funded by the European Union

restricted under its mis- and disinformation policies by EU Member States. In addition, LinkedIn provides information on misinformation content removed from its platform across EU member states for the reporting period, including how many of these removals were appealed, and how many appeals were granted (pp.28-29). This data indirectly addresses the question of how much disinformation can be found on the platform (SI1) and how much of it has been demonetized (SI4), however, without context. For instance, the report does not indicate a percentage of posts that are disinformation, or a share of content flagged as disinformation relative to the total of posts. It therefore remains difficult to assess to what extent that platform has adequately and proportionally responded to the threats of disinformation.

Instagram & Facebook (Meta)

The 139-page report by Meta includes a clear statement of its mis- and disinformation policy, namely that “Misinformation is considered to be unacceptable content” (Meta, 2025, p.16). The report confirms that Meta has signed up for Commitment 41 for both of its platforms Facebook and Instagram. Meta further agrees to new implementation measures, referring to “new tools and policies” developed to meet this commitment. Their specification is that “We continue to engage with the Taskforce Monitoring Working Group” (p.117). Meta, like LinkedIn, also lists the number of ads removed for violating its misinformation policy across EU Member States, for Facebook and Instagram combined, as well as the overall number of ads removed for any policy violation. The report references the Ad Library, where examples of political ad labelling can be found (notably, Meta stopped serving political ads in October 2025). In addition, the report includes a table showing the number of unique pieces of content removed for violating “harmful health misinformation or voter and census interference policies” in EU Member States. Meta also provides data on content viewed and treated with fact checks following assessments by third-party fact-checkers. These numbers give some insight into the prevalence of disinformation on the platform (SI1) and how much of it has been demonetized, e.g. when considering ad removals (SI4). However, without a denominator or baseline, these numbers are hard to interpret in terms of prevalence. Lastly, Meta states the number of partnerships with fact-checking organizations in Europe, thus giving insight into their platform investments in fact-checking (SI5). However, with a lack of financial metrics, it is difficult to reliably assess the scale of those actions.

TikTok

TikTok provides the longest report, spanning 308 pages. The platform indicates that it has not deployed new implementation measures under Commitment 41, specifying that it is “pending further updates from the Commission.” (TikTok, 2025, p.243). Notably, the platform also scaled back its commitments on fact-checking and empowering users, which are crucial service-level indicators that may explain some of the findings. Page 12 of the report contains a disclaimer that "TikTok remains committed to protecting our community from disinformation, strengthening resilience against misinformation, and upholding transparency and accountability throughout electoral and crisis contexts." Moreover, TikTok provides a variety of different and extensive quantitative data. This includes, again, numbers of removals of ads that violated their political content and misinformation policy. Beyond this, TikTok provides information about the audience of disinformation (SI3), by listing impressions of removed ads violating their misinformation policy. Numbers are provided for removed fake accounts and their followers (SI2), as well as information about videos that violated their misinformation policy (SI1) or were labelled "AI-generated." TikTok also mentions how many views these videos had prior to the removal. Overall, this data gives an idea of the prevalence of disinformation on the platform (SI1), where it may come from, e.g. fake accounts (SI2), how much of it TikTok removes (SI4), as well as how often people are exposed to it (SI3). Again, it lacks a baseline measure of prevalence to put these absolute numbers into perspective and interpret them meaningfully. Moreover, TikTok gives no insight into their investment in fact-checking.

YouTube (Google)

Google provides a 214-page report covering their services Google Search, YouTube, and Google Advertising. However, only YouTube as a social media platform is of interest in this white paper, as its scope is limited to social media platforms. On behalf of its related services, Google has signed up for Commitment 41, stating that it has been a participant in the working group dedicated to developing Structural Indicators and has supported research on this topic (Google, 2025, p.171). However, similar to other platforms, Google explicitly withdrew from all fact-checking commitments related to integrating third-party fact-checking into YouTube before the Code became binding. In addition, Google states that it "will continue to support the publication of Structural Indicators, and work towards further honing their methodology and scope" (p.172). Like the other platforms, Google provides quantitative data on ad removals across EU member states and appeals thereof. In addition, the report contains data on channels and videos that have been identified as violating their misinformation policy, including information on how many views videos received before their removal. As such, Google gives some insight into disinformation prevalence (SI1) and exposure towards it (SI3), and to what extent it is taken down, which indirectly addresses demonetization (SI4). However, contextual metrics are lacking again. Similar to the reports of other platforms, it is difficult to assess the extent to which YouTube has addressed disinformation threats.

Table 1. *Overview of information provided regarding structural indicators across platforms, based on September 2025 self-reports.*

Platform	Commitment 41	Metrics provided	SIs addressed
LinkedIn	Not subscribed	• Misinformation removal	SI1, SI4
Instagram, Facebook (Meta)	Subscribed (Engagement in working group)	• Misinformation removal • Fake account removal • Fact-checking partnerships	SI1, SI4, SI5
TikTok	Subscribed (pending further updates from the Commission)	• Misinformation removal • Misinformation impressions • Fake account removal • Fake account engagement • AI-generated content labeling • Misinformation moderation and outcomes	SI1, SI2, SI3, SI4
YouTube (Google)	Subscribed (Engagement in working group)	• Misinformation removal • Misinformation impressions • Fake account removal • AI-generated content labeling	SI1, SI3, SI4

Answering our first research question (*To what extent do platform signatories report on structural indicators in their most recent biannual transparency reports?*), we conclude that the information on structural indicators provided in the VLOP transparency reports remains rudimentary and limited, even though most have expressed their dedication and subscribed to Commitment 41 (except LinkedIn). Thus, it is difficult to draw conclusions regarding the first structural indicator, as the data provided does not estimate the relative prevalence of disinformation, or what kind of narratives it involves (SI1). Moreover, the reports do not adequately map the sources of disinformation (SI2), such as attribution to specific actors or information about scale by origin, making it difficult to assess in what contexts disinformation is spread. Furthermore, the reports do not provide details about audiences or reach beyond simplistic interaction counts (SI3), many of which are missing altogether. While platforms like Google and TikTok report removals of suspected content (e.g., videos or fake accounts), they only disclose information on demonetization indirectly, that is, when considering ads that have been restricted, but do not explicitly address financial penalty mechanisms (SI4). Most reports – except for Meta – offer no meaningful insight into fact-checking partnerships or investments (SI5). However, even Meta limits its reporting to coordination efforts and the disclosure of ongoing operations, without further financial or operational detail.

Our observations are in line with extant literature on structural indicators under the Code. Specifically, Nenadic et al. (2023) highlight in their development of structural indicators that metrics must be presented in context to adequately assess platform implementation – an action not yet undertaken. Research by Mündges and Park (2024) who conducted an extensive analysis of the January 2023 transparency reports also emphasize the striking lack of quantitative data. They conclude that “the commitment by signatories to develop and implement Structural Indicators remains unfulfilled” (p.15). This conclusion still holds. In addition, our analysis shows that the data provided by platforms is self-selected and non-standardized, with each platform reporting very different metrics. This allows platforms to appear compliant at a surface level, as the figures they submit may be drawn from the easiest or most accessible samples. However, without contextual information on prevalence, proportionality or details on the removed content, it is impossible to assess their true level of compliance with the code.

Monitoring via independent academic research

Platform self-reports are an important pillar of monitoring platform compliance, but it is not the only way to assess platform implementation of the *Code of Conduct on Disinformation*. Given that both our own analysis of recent transparency reports as well as previous investigations revealed that platforms themselves are providing insufficient information on structural indicators, we supplement our monitoring effort with independent scientific research. Specifically, we seek to understand to what extent we can learn about the implementation of the first three indicators, namely prevalence of disinformation (SI1), its sources and origins (SI2), as well as audiences (SI3) based on existing academic research. As academic literature mostly uses the term misinformation (rather than disinformation), given that intent is difficult to prove, we refer to misinformation in this section.

State-of-the-art literature review: The prevalence of misinformation on social media
A growing number of empirical studies in the academic literature have taken on the task to monitor, or at least make an estimate of, the amount of misinformation that people may come across when accessing different social media platforms. Thus, there are some estimates available in existing academic studies on the relative presence of

Co-funded by the European Union

misinformation across platforms. However, it is important to note that these studies were conducted over different timeframes than those covered in our analysis of the transparency reports, and are therefore not directly comparable. In many cases, the research focuses on specific topics or types of disinformation, whereas platform self-reports often do not distinguish in this way (except for health misinformation). Additionally, some studies do not differentiate between individual platforms, providing only general insights into disinformation exposure as a whole. For these reasons, we do not aim to draw direct comparisons between the scientific literature and our own analysis, but instead to highlight broader trends and synergies.

Regarding prevalence and exposure to misinformation overall, research finds that it might only constitute a small amount of all information citizens come across in their daily news diets. For instance, in a review paper of different approaches to access misinformation diets in data donation studies and web tracking, Acerbi et al. (2022) conclude that exposure to online misinformation overall ranges between 0.1% (Altay et al., 2022) and 6% (Guess et al., 2020). These estimates are based on domain-level measures of exposure to certain forms of online desktop or mobile media. Only Boberg et al. (2021) in this overview has relied on analysis at the level of posts, but defined misinformation in a restrictive way, namely as “conspiracies or fake news content.” Boberg et al.’s (2021) findings showed that German users of Facebook were exposed to these malicious contents via alternative news media about 1.1% of the time in the first 2.5 months of 2020. They also found that posts from alternative news channels generated less engagement than those from mainstream news channels.

Guess et al. (2019) examined the prevalence and predictors of fake news dissemination on Facebook using a combination of survey and behavioral data. Their research indicates that sharing fake news is a rare activity among the average American Facebook user. This is evidenced by the fact that the majority of Facebook users, both Democrats and Republicans, did not share any links to fake news content in the year 2016. Still, among the minority that did share at least one fake news post (8.5% of users), Republicans were overrepresented with 18.1% of Republicans who shared at least one fake news post, while among Democrats this was only 3.5%. This suggests that while the overall prevalence of fake news on Facebook is lower than commonly perceived, a small group of users is disproportionately responsible for its spread.

Yang et al. (2023) extended their analysis to visual content on Facebook analyzing more than 13.7 million posts from August to October 2020. Unlike other studies that focused on misinformation in text and at the domain level, this study focuses on visuals and finds a substantially higher proportion of misinformation. About 20% of images that depicted political figures were identified as misinformation, just as were 23% of images that depicted any political content relevant to the US. This appears high and may be explained by the fact that AI image generators were already widely available during data collection.

While some platforms, such as Twitter/X and Facebook, are somewhat well-studied when it comes to misinformation prevalence, the same cannot be said for more recent platforms, such as Instagram or TikTok. Research on these platforms has primarily focused on user perceptions when it comes to misinformation prevalence. In a recent study surveying citizens in the US, UK, France, Canada, and Germany about their perceptions of misinformation on Instagram, Hoffmann and Boulianne (2025) found that citizens say that they see it sometimes to moderately often. They also indicate to be moderately concerned about misinformation on Instagram. On TikTok, Kirkpatrick and Lawrie (2024) find somewhat more concerning results when studying perceptions of health misinformation among women. Almost all surveyed users (i.e., 98.15%) perceived health misinformation to be quite prevalent on TikTok and they considered the impact to be quite serious.

Limitations in academic research

While these findings provide valuable insights, we also want to point out some limitations and considerations that impede the generalizability and validity of extant research. These may have implications for policy and practice, as they mostly pertain to the nature of misinformation claims and the difficulty in identifying misinformation. We also reflect on the difficulty of discerning between mis- and disinformation in this section of the white paper.

Limits of domain-level analyses. First and foremost, existing data on the level of misinformation across platforms is mainly based on a domain-level estimate of exposure to misinformation (see Acerbi et al., 2022; Grinberg et al., 2019). Concretely this means that studies tend to use exposure not to specific content, but to sources that are identified to generally spread false, deceitful, conspiratorial or inaccurate information. The main issue with this crude approach is that consuming certain sources with an affinity for conspiracies, satire, or hyper-partisan views is equated with misinformation exposure, while little is known about the exact content that citizens saw. Crucially, misinformation exists on the content level, and not on the source level: One cannot assume that all information accessed through suspicious domains is misinformation (see Altay et al., 2023; Hameleers, 2025). This is also reflected in the operationalizations of misinformation used in many existing studies. Cordonier et al. (2021), as an example, referred to misinformation as sources with a high likelihood for spreading conspiracy theories, clickbait headlines, satire, and pseudo-science. Osmondsen et al. (2021) referred to information with little regard to the truth as misinformation. In this definition, it remains unclear whether all misinformation has little regard for the truth, and whether having little regard for the truth can define misinformation (i.e., opinion-related content may also fall under the scope of misinformation if one would apply this definition). These definitions do not align with the common definition of misinformation as verifiably false or inaccurate information on the level of content. The first issue thus mainly relates to the conflation of suspicious sources with an affinity for hyper-partisan or conspiracy views with misinformation on the content level.

Misinformation is more granular and nuanced than assumed. For our evaluation of the implementation of the *Code of Conduct on Disinformation* this means that measures of structural indicators 1 and 2 are lumped together. The second issue relates to the parameters of misinformation, and specifically the boundaries between true and false content drawn in existing studies. The issue is that misinformation does not exist as a simple boundary between true and false information – which would be directly applicable to the monitoring and auditing requirements of the code. Misinformation is highly granular, consists of different degrees of deception, and relies on various epistemologies that are either close to or very different from the ways in which true reporting is constructed (Hameleers & Yekta, 2025). This goes against existing work that has reduced misinformation and its identification to certain distinguishable linguistic and stylistic features in misinformation, such as negativity, emotionality, and the circumvention of expert sources (Carrasco-Farré, 2022; Damstra et al., 2021; Lebernegg et al., 2025).

Misinformation often relies on visuals and decontextualized evidence. A related issue is that research mapping the prevalence of misinformation mostly looks at clearly identifiable misinformation based on a restricted set of distinctive features. This approach may fail to capture subtle forms of misinformation, such as visual decontextualization or a re-interpretation of empirical evidence (see e.g., Brennen et al., 2021). They may also overlook forms of misinformation that are based on mimicking established news in terms of layout, content, or argumentative structure (Hameleers & van der Goot, 2024). Crucially, if we assume that the evaluation of the code requires data on the relative amount of misinformation as a proportion of all information present and promoted on distinct social media platforms, we need to be able to map different forms of misinformation and deviations from true content.

Although the overview presented here is not complete and plagued with potential validity and reliability issues, it can be concluded that misinformation may not make up the majority of all information, albeit its presence is salient across social media and across domains of information. Overall, the reviewed studies show that misinformation prevalence fluctuates over time and with different topics. They also note that misinformation is concentrated among specific audiences, such that low average prevalence masks the fact that for some audience segments the exposure of misinformation is substantially higher. In conclusion, we suggest the following as a way forward for future studies investigating the scope of mis- and disinformation on social media:

- **Analyze content directly** rather than assuming entire sources are disinformation and focusing on evaluating specific pieces of content.
- **Measure disinformation proportionally**, assessing how much disinformation appears relative to all information users encounter on a platform.
- **Differentiate types and formats of disinformation**, including modality, style, and subtle forms like decontextualization or partial falsehoods.
- **Expand detection beyond obvious cases**, ensuring subtle, fact-adjacent disinformation is identified, not just extreme or visible examples.
- **Avoid blind spots** by accounting for misleading content that appears credible due to its similarity to accurate information.

The limitations of research also come with implications for policy and the implementation of the code. For instance, the complexity of misinformation and the subtle manners of decontextualization may make regulation more difficult – as various practices may remain unnoticed or fall under freedom of speech. Another main issue that affects the issues mentioned above is that intentions are difficult to prove based on lacking details about motivations and contexts in which misinformation is created and spread. Thus, although looking at the context may help to identify intentions, the assessments of intent will be based on a probability score that may lack empirical proof.

Empirical case study: Prevalence of disinformation and common narratives in Dutch language regions (The Netherlands/Belgium)

Thus far, we have demonstrated the difficulty of researching structural indicators based on platform self-reports and identified common limitations of existing academic studies. To address these challenges, we present an original computational analysis of social media data covering the same period as the analysis of platform self-reports (January 2025–June 2025). The analysis is loosely inspired by previous investigations conducted by two research institutions commissioned by EDMO: TrustLab and Science Feedback. However, rather than examining multiple EU countries in parallel, our analysis focuses on Dutch-language data, thus encompassing content from the Netherlands and/or Belgium – the two countries represented by the BENEDMO hub. Moreover, this study focuses in depth on the first structural indicator (SI1) – the prevalence of disinformation – by uncovering and analyzing common disinformation narratives. Importantly, rather than relying on fact-checkers to annotate the data, our approach is fully automated, making the investigation quickly repeatable. By analyzing content directly and not assuming entire sources are disinformation, measuring the prevalence of disinformation relative to other forms of problematic content, and accounting for varying levels of granularity through

Co-funded by the European Union

in-depth narrative analysis, our approach overcomes several limitations identified in prior academic research. With this case study we aim to answer the questions:

RQ2: *How prevalent is disinformation on social media platforms in the Dutch-language context, based on data covering 1 January–30 June 2025?*
RQ3: *Which disinformation narratives emerge in this context?*

A look back: Previous EDMO-investigations by TrustLab and ScienceFeedback
In 2023, TrustLab (2023) conducted pioneering research on the prevalence (SI1) and sources of disinformation (SI2) in Poland, Slovakia, and Spain. At the time, X was still a platform signatory, hence it is included. This report shows disinformation discoverability – which the authors define as the ratio of mis/disinformation posts among sensitive content – is overall highest on Twitter, followed by Meta's services Facebook and Instagram. Discoverability of disinformation was found to be lowest on YouTube. Moreover, the authors show that mis -and disinformation content showed higher engagement than non-mis/disinformation content on Twitter, while the opposite was true for TikTok. In addition, a recent analysis by ScienceFeedback (Vincent et al., 2025) gives insights into several structural indicators (SI1, SI2, SI4) of disinformation on platforms in France, Poland, Slovakia and Spain. Somewhat contradicting the TrustLab findings, Science Feedback shows that TikTok has the highest prevalence of misinformation across the four EU-countries, (cira 20% of exposure-weighted posts on public-interest topics), followed by Facebook (13%) and X/Twitter (11%). Instagram and YouTube are lower (8%), and LinkedIn was found to have the lowest (2%).

BENEDMO case study
Our analysis is based on the social media monitoring tool EHBT.be (in Dutch: *"Eerste Hulp bij Twijfel"*; in English: *"First Aid When in Doubt"*), which is designed to track prevalent social media narratives. On the platform, users of the dashboard can create custom channels by entering a list of keywords and a language, after which automated collectors will search for occurrences of those words on a variety of platforms - including current Code signatories Facebook, Instagram, TikTok and YouTube, but also 4chan, 9gag, Reddit, Telegram, Threads, Twitter and a collection of alternative news media, (mostly) Wordpress sites that are known distributors of disinformation. We did not have access to data from LinkedIn, which is why this signatory is not included in our analysis.

Co-funded by the European Union

After anonymization, machine learning classifiers were used to provide metadata-layers, such as a toxicity-score, a score representing whether the user expressed a question for information and a score representing the factuality of the statements made. Every day, the new discussion threads were summarized using Google Gemini 2.5. To prevent hallucinations the minimum number of items in a discussion needs to be 10 and an additional LLM-pass is done for each summary to double check whether the summary contains only information from the discussion thread itself.
For our analysis, we used data from 20 EHBT-channels. These range from topics that are highly vulnerable to disinformation to topics that typically elicit heated discussions and abusive rhetoric. Specifically, these are: vaccination and COVID; Israel-Palestine conflict; refugees and social security; Russia; foreign interference in elections; 5G; whooping cough; transgender issues; chemtrails; Sahara sand; bee mortality; monkeypox; H.A.A.R.P.; food safety; anorexia; contraception; self-help and self-diagnosis; cosmetic procedures; supplements; hormone therapy; weight loss; climate change (see also Figure 1). In the Appendix, we provide an overview of the exact keywords used, the channels this study was based on, including the search terms that EHBT monitored for. In total, we aggregated 72k (unique) posts and 5M comments from this analysis. To identify periods of unusual activity in our dataset, we employed complementary anomaly detection techniques, which are also reported in the Appendix.

Method:
We extracted approximately **72,000 unique posts and 5 million comments** from a wide range of social media platforms in the Dutch language, based on keyword searches covering topics highly vulnerable to disinformation as well as topics that typically elicit heated discussions and abusive rhetoric. Content classification was conducted in **three iterations using large language models**: Google Gemini 3 Pro and OpenAI's GPT-4.1 independently annotated social media discussions, achieving substantial agreement (Cohen's Kappa = 0.72), with Claude Opus 4.5 acting as a tiebreaker for disputed cases.

Annotation
Our data analysis was inspired by the Science Feedback Report "Measuring the State of Online Disinformation in Europe on Very Large Online Platforms" (Vincent et al., 2025), in an attempt to make the two reports somewhat comparable. The authors define a number of categories, of which we retained the following subset relevant to our research goals:

- **Mis/disinformation**: Content stating or clearly implying a verifiably false or misleading claim that may cause public harm.
- **Borderline**: Content feeding a misleading narrative without necessarily containing outright falsehoods, but potentially reinforcing false beliefs.

Co-funded by the European Union

- **Abusive**: Content not containing mis/disinformation but involving harmful material such as hate speech, insults, spam, or incitement to harmful behaviour.
- **Unverifiable**: Content that cannot be assessed as either credible or mis/disinformation (e.g. opinion-based).
- **Credible and informative**: Content conveying true or credible information on important matters about the state of the world (excluding trivia, gossip, or anecdotes).
- **Irrelevant**: Content not about public affairs or scientific/political issues (e.g. entertainment, sports, religious content, cooking recipes without health claims, geographically irrelevant to Europe).

Since expansive manual data annotation was out of scope for this white paper, we built an automated pipeline that leverages recent advances in the use of LLMs-as-a-judge. Vincent et al. (2025) outline some preliminary experiments using LLMs for annotation and observed an accuracy between 50% and 70%, although it should be noted that fairly weak models were used. We attempted to improve on the pipeline, by using several state-of-the-art LLMs correcting each other.

Iterations

Iteration 1: Google Gemini 3 Pro (1M context) was presented with each social media post, the associated comments and the summary of the discussion. It was then asked to label the content with 1 of the aforementioned categories, as described above. Note an important difference with the Vincent et al. (2025)-approach which operates on the level of the post, whereas our analysis concentrates on the discussion. This means we need to interpret for example the label "Mis/Disinformation" as: the social media post contained mis/disinformation **or** elicited many comments containing mis/disinformation. This has the advantage of capturing the prevalence of problematic content in the comment feeds of reputable news sources as well.
Iteration 2: In the next iteration, we used OpenAI's GPT-4.1 (1M context) to perform the same task, rendering a parallel annotation for each post. Cohen's Kappa-score for these annotators was fairly high at 0.72.
Iteration 3: Finally, we used Claude Opus 4.5 to be a tiebreaker for the roughly 15k posts for which there was disagreement. Claude was presented only with the two options from the first iteration to choose from.
While we acknowledge the limitations of the automated approach, it has the distinct advantage of being able to produce reliable results at the touch of a button, enabling the automated generation of these types of analyses for very specific time frames and fringe topics.

We performed a manual spot check of 80 items for each of the 6 labels to estimate the accuracy of the automated annotations. We observed an accuracy of 78.8% over all labels, with the most confusable category combinations being Borderline vs Mis/Disinformation, Borderline vs Unverifiable, and Unverifiable vs Irrelevant. 1 edge case had the automated approach label credible information as misinformation, when the death of a holocaust denier was discussed, although the discussion did not amplify holocaust denial itself. When looking at the aggregate categories of what Vincent et al. (2025) call problematic content (abusive/borderline/mis/disinformation) vs other (credible & unverifiable), we observed an encouraging accuracy of 93.6%.

Results

Answering RQ2 (*How prevalent is disinformation on social media platforms in the Dutch-language context?)*, our analysis suggests that topics often considered fringe, such as 5G, H.A.A.R.P., and chemtrails, are associated with relatively higher levels of mis/disinformation and borderline content. Medically related topics, including anorexia and dietary supplements, appear to be less affected overall; however, for broader and more prominent thematic channels – such as vaccination and COVID-19, climate change, and the Israel–Palestine conflict content classified as problematic (i.e. mis/disinformation, borderline, or abusive) occurs in at least half of the discussions observed in our dataset (see Figure 2).

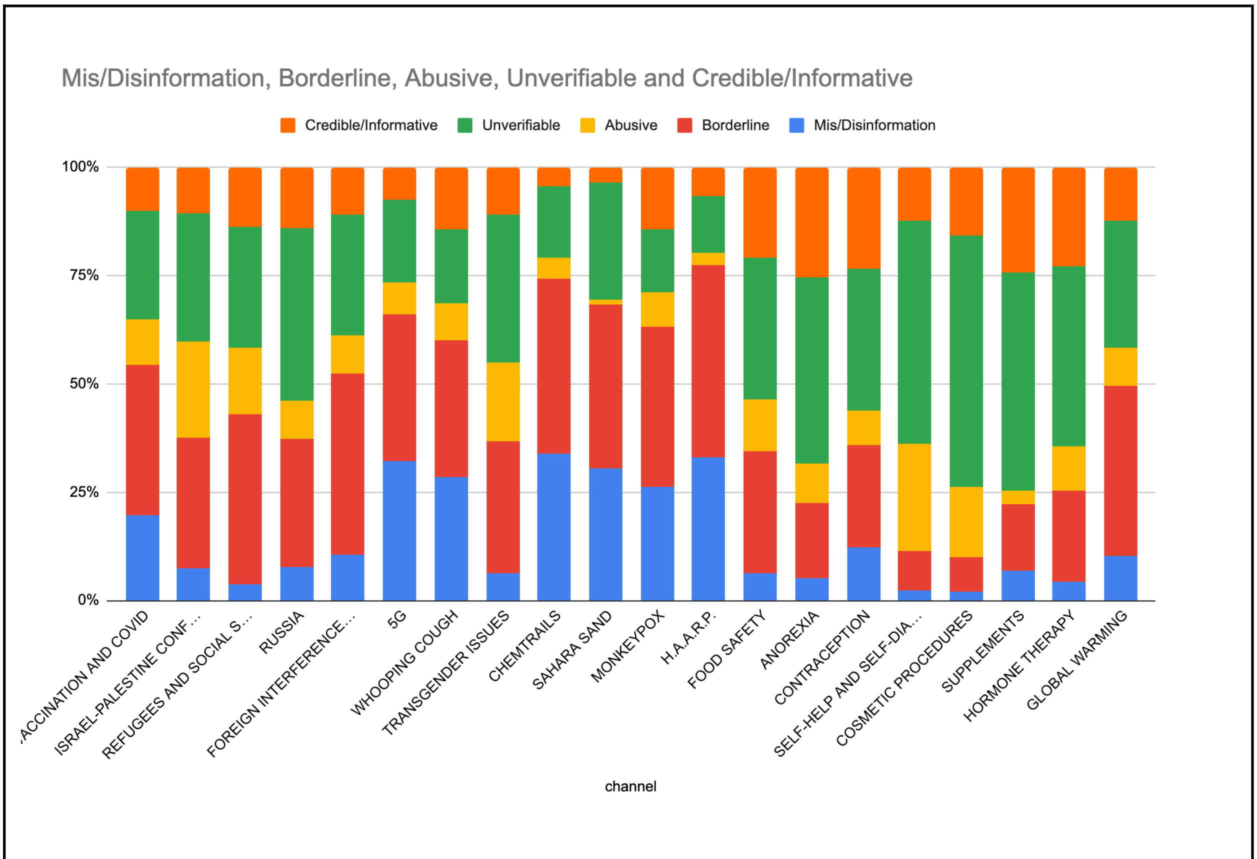

Figure 2. *Labels per topic (=EHBT-channel).*

When examining label distribution by platform (see Figure 3), we observe that among the Code signatories, Instagram exhibits the highest percentage of clear-cut mis-/disinformation (5.8%), followed by Facebook (4.9%), TikTok and YouTube (each 2.4%). When including borderline content – i.e., content that feeds a misleading narrative without necessarily containing outright falsehoods but may reinforce false beliefs – Facebook leads this ranking with 29.6% followed by Instagram (21.5%), TikTok (23.3%), and YouTube (19.5%).

In general, it seems like there is less disinformation on platforms who are signatories than on those who have not signed the Code. Regarding the latter, however, it should be noted that the 4chan and 9GAG columns represent very limited data, as Dutch-language content is rare. Problematic content on Telegram appears very high, although this may partly reflect the types of channels and groups monitored by EHBT. The aggregated WordPress category (covering sites such as NineForNews and FrontNieuws) also shows a high proportion of problematic content, with only Twitter and Telegram exceeding it. Finally, Reddit – often characterized as a toxic platform – shows only 25% problematic content in total (abusive, borderline, disinformation), though Dutch-language posts are limited, and moderation in subreddits may contribute to this lower figure.

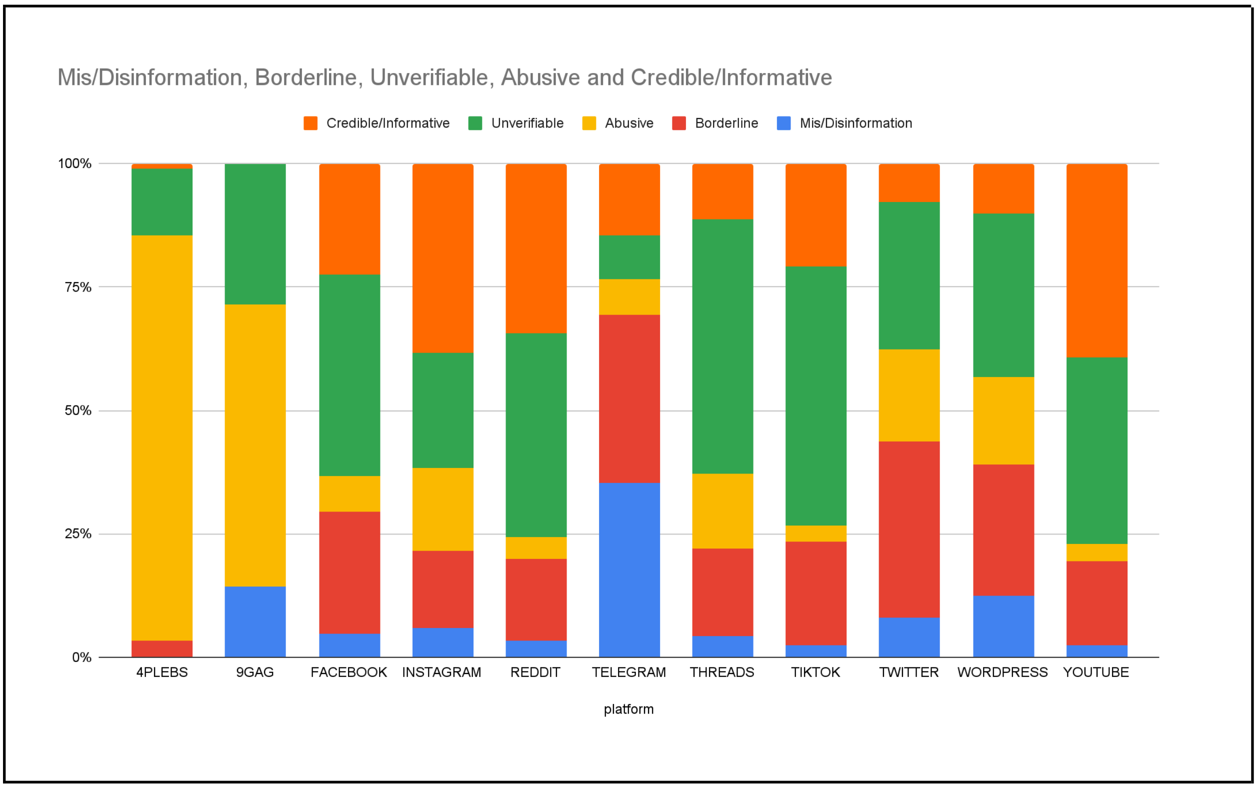

Figure 3. *Labels per platform.*

To answer our third research question (*Which disinformation narratives emerge in this context?*) we zoomed in on the rhetorical patterns and narrative structures present in Dutch-language online discourse, collected by the EHBT-social media monitor in the channel “Refugees and Social Security.” The analysis is based on 7,465 flagged social media posts and their associated comments. Table 2 below gives an overview of some of the more prominent disinformation narratives observed, with citations to illustrative examples. The narrative taxonomy presented here documents six primary clusters: demographic replacement, institutional conspiracy, social breakdown, geopolitical conflict, health/science denial, and resource competition. These clusters exhibit interconnection, with individual posts frequently combining elements from multiple categories, for instance, linking “omvolking” (the “Great Replacement” theory) to World Economic Forum agendas, or connecting vaccine narratives to “Great Reset” frameworks. The rhetorical patterns recur across platforms, though Telegram hosts a higher concentration of explicit claims while mainstream platforms more frequently contain content that implies rather than states conclusions directly.

Table 2. *Clusters of prominent disinformation narratives in the Dutch context.*

Overall narratives	**Sub-narratives**	**Explanation**
Demographic Replacement Narratives	• “Omvolking” (Great Replacement) • The Kalergi Plan • “Fifth Column” and “Enemy Within” Rhetoric	The dataset contains 623 posts referencing *“omvolking”*, the Dutch term for the “Great Replacement” theory, framing immigration as an existential threat to national identity. Related conspiratorial narratives include 51 posts citing the “Kalergi Plan” - a conspiracy theory claiming European elites are engineering ethnic replacement - and around 10 posts portraying immigrants as a “fifth column” or internal threat, with additional content using broader “enemy within” framing.
Institutional Conspiracy Narratives	• World Economic Forum and “Great Reset” • “Deep State” Narratives • Media as Propaganda Apparatus	This cluster portrays national and supranational institutions as engaged in coordinated conspiracies against citizens. The term *“deep state”* appears in 53 posts, depicting governments as facades for hidden power, while *“propaganda”* appears in 258 posts, often targeting mainstream media, particularly the Dutch public broadcaster NOS, which is framed as evidence of coordinated deception.

Social Breakdown and Violence Narratives	• "Burgeroorlog" (Civil War) Rhetoric • Islamization and Sharia Narratives	The term *"burgeroorlog"* (civil war) appears in 43 posts, framing social tensions as potential armed conflict and often calling for exceptional measures. References to *"sharia"* appear in 42 posts, portraying Muslim presence as a threat to national law.
Geopolitical Conflict Narratives	• Israel-Gaza war • Russia-Ukraine war	Posts about the war in Gaza regularly present contested characterizations of military operations as established fact, with parallel narratives disputing casualty figures and source credibility. Narratives concerning the war against Ukraine include claims about fabricated threats, corruption characterizations, and framing of Western support as elite manipulation.
Resource Competition Narratives	• Housing Priority Claims • Financial Burden Claims	This narrative cluster frames immigration through resource scarcity, presenting asylum seekers as receiving preferential access to housing, benefits, and services at citizens' expense.

Our case study analysis has several limitations, some of which we have also identified in the existing academic literature. First, our analysis starts from a predefined set of topics that are particularly prone to disinformation, which limits our ability to extend the analysis beyond more obvious cases. Moreover, as it is not feasible to obtain a complete overview of all content circulating on social media platforms at any given moment, we cannot provide definitive estimates of the relative prevalence of disinformation across platforms. However, our approach does allow us to assess the prevalence of disinformation relative to other forms of problematic content, such as abusive or borderline content, and to examine how much of the information within disinformation-prone topic areas is accurate. We also demonstrate that disinformation is not always clear-cut but often exists along a continuum, which is captured by our "borderline" category. Finally, our taxonomy of narrative clusters provides a more detailed understanding of the types of disinformation and problematic discourse circulating in the Dutch-language context. This, in turn, can support political actors and media practitioners in addressing emerging narratives and areas of concern.

Conclusion

In this white paper, we address the challenge of monitoring the implementation of the *Code of Conduct on Disinformation* across social media platforms through three complementary approaches: (1) an assessment of platform signatories' self-reports, (2) a review of existing academic literature on the prevalence of disinformation on platforms, and (3) a case study examining the prevalence of disinformation and common narratives in Dutch-language regions.

Overall, our findings indicate that platforms disclose limited information in their self-reports. These reports often lack comparable and transparent metrics and provide little insight into the specific types of problematic content that have been removed, making it difficult to assess actual compliance with the Code. Both the academic literature and our case study suggest that the overall prevalence of mis- and disinformation is relatively low, typically accounting for only a small percentage of content. However, they also highlight that mis- and disinformation are rarely clear-cut categories. While overt conspiracy theories are often easier to identify, more ambiguous or borderline content –which may nonetheless reinforce false beliefs – remains difficult to detect and assess. Given the granular nature of such content, understanding the full scope of the issue remains challenging. Our analysis further shows that, despite platforms' commitments under the Code, disinformation continues to circulate, pointing to the inherent limitations of existing platform measures.

Based on these findings, we recommend that platforms provide greater transparency regarding the prevalence of mis- and disinformation on services, and ideally also the types of problematic content they find and remove. This would support the development of more targeted mitigation strategies, for example by fact-checkers and media literacy practitioners. From a policy perspective, **incorporating a narrative dimension into the set of structural indicators could therefore be beneficial**, as it would encourage platforms to move beyond self-selected metrics and engage more directly with the substantive content they identify as problematic.

Disclaimer

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HADEA). Neither the European Union nor the granting authority can be held responsible for them.

BENEDMO

References

Acerbi, A., Altay, S., & Mercier, H. (2022). Research note: Fighting misinformation or fighting for information? *Harvard Kennedy School Misinformation Review*. https://doi.org/10.37016/mr-2020-87

Altay, S., Berriche, M., & Acerbi, A. (2023). Misinformation on Misinformation: Conceptual and Methodological Challenges. *Social Media + Society*, *9*(1), 205630512211504. https://doi.org/10.1177/20563051221150412

Altay, S., Kleis Nielsen, R., & Fletcher, R. (2022). Quantifying the "infodemic": People turned to trustworthy news outlets during the 2020 coronavirus pandemic. *Journal of Quantitative Description: Digital Media*, *2*. https://doi.org/10.51685/jqd.2022.020

Bleyer-Simon, K., & Reviglio, U. (2024). *Defining Disinformation across EU and VLOP Policies*. https://edmo.eu/wp-content/uploads/2024/10/EDMO-Report-%E2%80%93-Defining-Disinformation-across-EU-and-VLOP-Policies.pdf

Boberg, S., Quandt, T., Schatto-Eckrodt, T., & Frischlich, L. (2020). Pandemic populism: Facebook pages of alternative news media and the corona crisis--A computational content analysis. *arXiv preprint arXiv:2004.02566*.

Böswald, L.-M. (2025). *Soft Law, Hard Risks? Co-Regulation and Risk Mitigation Under the Digital Services Act*. https://www.interface-eu.org/publications/dsa-co-regulatory-mechanisms

Carrasco-Farré, C. (2022). The fingerprints of misinformation: how deceptive content differs from reliable sources in terms of cognitive effort and appeal to emotions. *Humanities and Social Sciences Communications*, *9*(1), 162. https://doi.org/10.1057/s41599-022-01174-9

Cordonier, L., Cafiero, F., & Bronner, G. (2021). Why are conspiracy theories more successful in some countries than in others? An exploratory study on Internet users from 22 Western and non-Western countries. *Social Science Information*, *60*(3), 436–456. https://doi.org/10.1177/05390184211018961

Damstra, A., Boomgaarden, H. G., Broda, E., Lindgren, E., Strömbäck, J., Tsfati, Y., & Vliegenthart, R. (2021). What Does Fake Look Like? A Review of the Literature on Intentional Deception in the News and on Social Media. *Journalism Studies*, *22*(14), 1947–1963. https://doi.org/10.1080/1461670X.2021.1979423

Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election. *Science*, *363*(6425), 374–378. https://doi.org/10.1126/science.aau2706

Guess, A. M., Nyhan, B., & Reifler, J. (2020). Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behaviour*, *4*(5), 472–480. https://doi.org/10.1038/s41562-020-0833-x

Co-funded by the European Union

Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, *5*(1). https://doi.org/10.1126/sciadv.aau4586

Hameleers, M. (2025). The visual nature of information warfare: the construction of partisan claims on truth and evidence in the context of wars in Ukraine and Israel/Palestine. *Journal of Communication*, *75*(2), 90–100. https://doi.org/10.1093/joc/jqae045

Hameleers, M., & van der Goot, E. (2025). Look at what the real facts and experts say! The use of expert references and objectivity claims in disinformation: A qualitative exploration and typology. *Journalism*, *26*(7), 1469–1487. https://doi.org/10.1177/14648849241257383

Hameleers, M., & Yekta, N. (2025). Entering an Information Era of Parallel Truths? A Qualitative Analysis of Legitimizing and De-legitimizing Truth Claims in Established Versus Alternative Media Outlets. *Communication Research*, *52*(6), 803–825. https://doi.org/10.1177/00936502231189685

Hoffmann, C. P., & Boulianne, S. (2025). Concerns about misinformation on Instagram in five countries. *Journal of Elections, Public Opinion and Parties*, *35*(3), 409–431. https://doi.org/10.1080/17457289.2025.2514192

Kirkpatrick, C. E., & Lawrie, L. L. (2024). TikTok as a Source of Health Information and Misinformation for Young Women in the United States: Survey Study. *JMIR Infodemiology*, *4*, e54663. https://doi.org/10.2196/54663

Lebernegg, N., Eberl, J.-M., Tolochko, P., & Boomgaarden, H. (2025). Do You Speak Disinformation? Computational Detection of Deceptive News-Like Content Using Linguistic and Stylistic Features. *Digital Journalism*, *13*(8), 1373–1398. https://doi.org/10.1080/21670811.2024.2305792

Li, H. O.-Y., Pastukhova, E., Brandts-Longtin, O., Tan, M. G., & Kirchhof, M. G. (2022). YouTube as a source of misinformation on COVID-19 vaccination: a systematic analysis. *BMJ Global Health*, *7*(3), e008334. https://doi.org/10.1136/bmjgh-2021-008334

Mündges, S., & Park, K. (2024). But did they really? Platforms' compliance with the Code of Practice on Disinformation in review. *Internet Policy Review*, *13*(3). https://doi.org/10.14763/2024.3.1786

Nenadic, I., Brogi, E., & Bleyer-Simon, K. (2023). *Structural indicators to assess effectiveness of the EU's Code of Practice on Disinformation*.

Nenadic, I., Brogi, E., Bleyer-Simon, K., & Reviglio, U. (2024). *Structural Indicators of the Code of Practice on Disinformation: The 2nd EDMO report*.

Osmundsen, M., Bor, A., Vahlstrup, P. B., Bechmann, A., & Petersen, M. B. (2021). Partisan Polarization Is the Primary Psychological Motivation behind Political Fake News

Sharing on Twitter. *American Political Science Review*, *115*(3), 999–1015. https://doi.org/10.1017/S0003055421000290

TrustLab. (2023). *A Comparative Analysis of the Prevalence of Misinformation and Sources of Disinformation across Major Social Media Platforms in Poland, Slovakia, Spain, and France*. https://www.trustlab.com/resources/codeofpractice-disinformation

Vincent, E., Carniel, B., & Crisan, D. (2025). *Measuring the State of Online Disinformation in Europe on Very Large Online Platforms First report of the SIMODS project (Structural Indicators to Monitor Online Disinformation Scientifically)* . https://science.feedback.org/first-measurement-disinformation-major-platforms-europe/

Yang, Y., Davis, T., & Hindman, M. (2023). Visual misinformation on Facebook. *Journal of Communication*. https://doi.org/10.1093/joc/jqac051

Appendix

Anomalies:

To identify periods of unusual activity in our dataset, we employed complementary anomaly detection techniques that distinguish between two fundamentally different phenomena: **volume spikes** (days with abnormally high posting activity) and **distribution shifts** (days where the proportion of content categories deviates from typical patterns).

Methodology

We aggregated the five classifications (*Mis/Disinformation*, *Credible/Informative*, *Unverifiable*, *Borderline*, and *Abusive)* by day to construct daily time series for each category. We then used Isolation Forest, an unsupervised machine learning algorithm to detect anomalies by measuring how easily a data point can be isolated through random partitioning. Unlike distance-based methods, Isolation Forest exploits the property that anomalous points require fewer random splits to isolate. We applied this algorithm in two configurations: once on raw daily counts (detecting volume anomalies) and once on daily proportions (detecting distributional anomalies).

The proportion-based approach is particularly relevant for misinformation research, as it is able to reveal days where the *composition* of content was unusual even when total posting volume was within normal bounds.

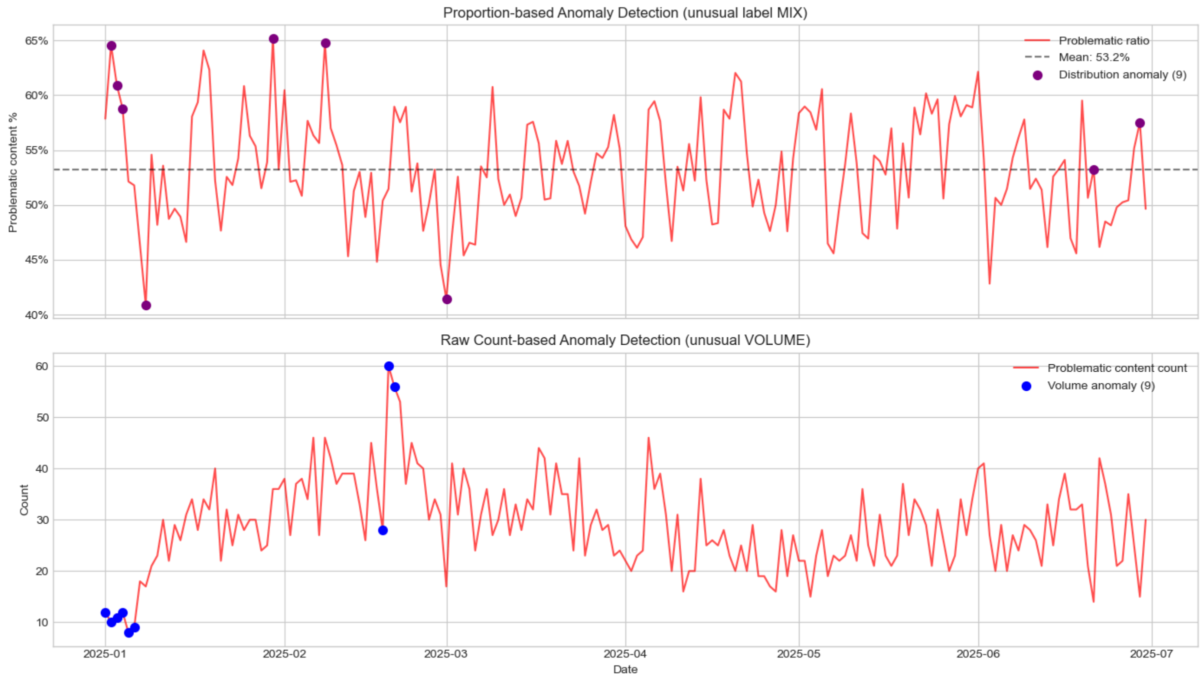

Figure: Problematic content anomalies

This figure displays the anomalies for problematic content (Mis/Disinformation, Abusive, Borderline categories). The volume spikes in February represent elevated activity in the Russia channel, instigated by renewed peace talks initiated by Donald Trump. However, this volume spike does not coincide with a distribution shift. There are very few distribution shifts of note. The early January ones are related to the lower total volume, causing small shifts to appear significant.

Notable distribution shifts can be linked to current events:

- January 30th marks an anomaly, which is due to the Abusive category (Mis/Disinformation and Borderline are within normal range). Inspection of the data shows the content mostly relates to the freeing of hostages by Hamas.
- The spike on February 8 represents the fallout after Donald Trump's claims that USAID paid news outlets for positive coverage of Democrats. The narrative spreads to other categories that link media control to conspiracies related to vaccination and coverage of the Israel-Palestine conflict
- Finally, June 30th is marked by an unusual distribution shift in the Borderline category. Discussions deal with the Glastonbury Festival controversy (Bob Vylan's "Death to IDF" chants), and Dutch musician Douwe Bob refusing to perform at a Jewish football event over "Zionist posters,".

Topics and search terms of EHBT-channels:

Topic	**Search Terms**
vaccination and COVID	anti-vaccinatie, astrazeneca, bijwerkingen vaccin, corona, corona maatregelen, coronaherstelfonds, coronaminister, coronastrategie, coronavaccin, coronavirus, coronavirussen, covid, covid, covid vaccinatie ervaring, covidvaccin, gevaccineerd, immuniteit, immuniteit, inenting, prik, prik, vaccin, vaccinatie, vaccinatie, vaccinatie en werk, vaccinatie kinderen, vaccinatie ouderen, vaccinatie verplicht, vaccinatie weigeren, vaccinatie zwanger, vaccinatiebewijs, vaccinatiecampagne, vaccinatiedebat, vaccinatiedecision, vaccinatiegegevens, vaccinatiegenocide, vaccinatiegraad,

Co-funded by the European Union

	vaccinatieinzicht, vaccinatiekritiek, vaccinatiemandaten, vaccinatieontwikkeling, vaccinatieonzin, vaccinatiepaspoort, vaccinatieplicht, vaccinatieregister, vaccinatieschade, vaccinexperts, virusverspreiding, wappies, zuigelingenvaccinatie, AstraZeneca, Big Pharma, COVID-19, HPV, Johnson & Johnson, MRNA, Moderna, Pfizer, bijwerkingen vaccin, boostershot, inenting, shedding, vaccin, vaccinangst, vaccinatieschade
Israel-Palestine conflict	Israel, antisemieten, antisemitische, antisemitisme, apartheidisrael, gaza, hamas, israel, israxl, israëlische, khamxs, palestijn, palestijnen, palestijns, palestijnse, palestinisering, yisrael, zionisme, zionists, ısrael
refugees and social security	asielbeleid, asielprocedure, asielzoekers, bijstand, immigratie, immigratieproblematiek, integratie, migratie, sociaal beleid, sociale dienst, sociale voorzieningen, sociale zekerheid, uitkeringen, uitkeringsfraude, vluchtelingen, vluchtelingen en werk, vluchtelingenstatus, vluchtelingenwerk
Russia	Krim, Oekraine, Putin, Rusland, Russisch, moscow, moskou
foreign interference in elections	buitenlandse inmenging, desinformatiecampagnes, informatieoorlog, inmenging in verkiezingen, kremlininvloed, verkiezingsbeïnvloeding, verkiezingsfraude, verkiezingsinmenging, verkiezingsmanipulatie
5G	5g, elektriciteitsoorlog

Co-funded by the European Union

whooping cough	kinkhoest
transgender issues	trans, transgender, transgenderacceptatie, transgenderactivisme, transgenderbewustzijn, transgenderdiscriminatie, transgenderexperiences, transgendergeschiedenis, transgendergezondheid, transgenderidentiteit, transgendermedischezorg, transgendermensrechten, transgendermilieu, transgenderminderheden, transgenderopvoeding, transgenderrechten, transgenderverschijnselen, transgendervraagstukken, transgenderzorg
chemtrails	chemtrail, chemtrails
Sahara sand	Sahara zand, saharazand
bee mortality	bijensterfte
monkeypox	apen pokken, apenpokken, mensapenpokken, moneypox, monkeypox, mpox, pokken
H.A.A.R.P.	H.A.A.R.P., HAARP
food safety	auditing, veiligheid, voedsel
anorexia	anorexia, braken, caloriebeperking, eetbuien, eetstoornis, gewichtsverlies, laxeren, lichaamsbeeld, ondergewicht, ondervoeding, purgeren
Contraception	artsen, bijwerkingen anticonceptie, condoom, cyclustracking, de pil, hormonenspiraal, hormoonvrij, implantaat, koperspiraal, menstruatie, menstruatiecyclus, morning-after, ovulatie,

	spiraal, vruchtbaarheid, vruchtbaarheidsbewustzijn
Self-help and self-diagnosis	Dr. Google, WebMD, adhd checker, adhd test, anxiety checker, anxiety test, autisme checker, autisme test, depressie checker, depressie test, depressietest, narcisme, selfchecker , symptomchecker, symptoom test, symptoomchecker, zelfchecker, zelfdiagnose, zelfhulp, zelfzorg
Cosmetic procedures	BBL, Brazilian Butt Lift, borstvergroting, borstverkleining, botox, buikwandcorrectie, chemische peeling, cosmetische ingreep, facelift, fillers, huidverjonging, huidverstrakking, injectables, kaaklijn, kinvergroting, laserbehandeling, lip lift, liposuctie, microneedling, mini-facelift, neuscorrectie, ooglidcorrectie, permanente make-up, plastische chirurgie, vetbevriezing, vetverwijdering
Supplements	AG1, BCAA, CBD, aminozuren, appelciderazijn, ashwagandha, collageen, collagen, creatine, eiwitpoeder, electrolyten, ginko biloba, magnesium, melatonine, multivitamine, omega 3, prebiotica, probiotica, superfoods, visolie, vitamines, voedingssupplement
Hormone therapy	HRT, hormonen, hormoonbalans, hormooncyclus, menopauze, oestrogeen, oestrogeentherapie, progesteron, puberteitsremmer, testosteron, testosteronboost
weight loss	afvallen, gewicht, gewichtsverlies
Climate change	klimaat, klimaatverandering