

Are Tech Companies Meeting Obligations to Protect Elections?

Evidence from Ireland's Presidential Election

August 2026



Authors

Aidan O'Brien: research and policy lead at EDMO Ireland.

Dr Eileen Culloty: coordinator EDMO Ireland, associate professor at DCU.

Acknowledgements

Monitoring online platforms during an election is a considerable undertaking, given the volume of content involved. The authors wish to thank the specialist researchers, civil society organisations, and partners from EDMO Ireland and across the EDMO network who provided assistance, tools, and data insights. We also acknowledge the platform representatives who responded to queries and made reporting channels available, and the European Commission, which has facilitated a Rapid Response System for elections across EU countries. Much has been learned through this process, and we hope this report contributes to constructive debate on how best to protect democratic elections.

About EDMO and EDMO Ireland

The European Digital Media Observatory (EDMO) was established in 2020 as part of an EU-wide effort to promote research on disinformation; advance fact-checking; support media literacy; and assess the implementation of the EU Code of Conduct on Disinformation.



The EDMO Network comprises 14 hubs operating across 27 countries in the EU and EEA. While hubs are co-funded by the EU and committed to common principles, they operate with full autonomy at national and regional levels.

EDMO Ireland includes the DCU Institute for Media, Democracy, and Society; TheJournal FactCheck; and NewsWhip. See <https://edmohub.ie/>

The report was updated on 26 August to clarify that the Rapid Response System is established by Signatories of the Code of Conduct on Disinformation

Summary

During elections in EU member states, Signatories of the Code of Conduct on Disinformation establish a Rapid Response System for communication between online platforms and actors monitoring disinformation and manipulation. The European Commission facilitates the process as Chair of the Code's Taskforce.

The regional hubs of the European Digital Media Observatory (EDMO) are invited to flag election-related content to platforms for review by participating platforms: Microsoft (LinkedIn), Meta (Instagram and Facebook), Alphabet/Google (YouTube) and TikTok. Platform participation is voluntary, but all platforms signatories have committed to participate under the Code. They are not obliged to do anything with the flagged content.

For the Irish Presidential Election in October 2025, EDMO Ireland monitored relevant platforms to identify content that appeared to be in breach of platforms' stated terms of services. This content was reported during the five week period of the Rapid Response System (24 September - 04 November). During this time:

- 177 items (pieces of content or accounts) were flagged to platforms. These were broken down as follows: election-related disinformation (85 items); AI-generated election content (26 items); harassment and abuse (13 items); incitement to violence (2 items); illegal election-day content (49 items); and issues relating to political advertising (2 items).
- Reported items included false claims about the legality of the nomination process; fake quotes attributed to public officials; assertions of systemic election fraud in Irish elections; AI-generated impersonations of candidates; illegal representations of spoiled votes; and videos of the harassment of political figures and canvassers.
- LinkedIn had the lowest number of reported items (7) while Facebook had the highest (54). Overall, platforms acted on half of the reported items. Responses by platform varied. Instagram was the most likely to act, responding to 59 percent of reported items, followed by TikTok (54%), YouTube (48%), and Facebook (44%). LinkedIn was the least likely to act, taking no action on reported items.
- Responses varied significantly by content type and platform. Platforms removed or geolocked all illegal content, but only acted on half of the AI-generated content. Only YouTube and TikTok acted on AI-generated content.

- Overall, platforms only acted on 13 of the 85 items flagged as electoral disinformation. Only Instagram and TikTok took action on these items.
- Of the 13 items reported as harassment and abuse, platforms only acted on 9 items. These included videos of public figures and canvassers being abused and threatened on the street.
- Reported items were generally addressed within 24 hours, but many reported items remained active for days, which is surprising given the urgency of an election. TikTok had the fastest response time to flagged content (averaging 6 hours).
- In many cases, the flagged content was associated with accounts that were monetised or had previously been banned. In addition, some items were posted by figures who had already come to public attention for propagating false claims or engaging in harassment.
- Overall, platform representatives often appeared unfamiliar with the context of the unfolding election and there was little evidence of pro-active measures. While items were reported as a potential breach of platforms' own terms of services, the platforms were repeatedly unable to explain inconsistencies in the application of their own rules. Signposting to authoritative sources, corrective information and other proactive information integrity efforts was poor across all platforms.

Introduction & Context

Digital platforms have become central to democratic elections, shaping campaign dynamics and the flow of information about candidates, issues, and the electoral process. In response to growing concerns about disinformation and electoral interference, major technology companies made public commitments to protect the integrity of elections.

Regulators and policymakers, at both national and EU levels, expect platforms to demonstrate that these commitments translate into effective and consistent implementation. This expectation is embedded in the EU's Digital Services Act, which obliges platforms to address systemic risks including risks to civic discourse and electoral processes¹. Platforms are expected to engage in proactive measures to address the amplification and monetisation of disinformation and manipulative behaviour; to take account of the specific circumstances of each election; and to provide meaningful transparency regarding their processes.

Platforms and Elections in Europe

Ireland's presidential election in October 2025 was the third major electoral event in Ireland since the implementation of the Digital Services Act. In the intervening period, online platforms received specific guidance from the European Commission² and Ireland's Electoral Commission³ on expectations around risk mitigation and responsiveness during elections.

Adopted in 2023, the Digital Services Act established a co-regulatory approach to the governance of very large online platforms. Under this model, the EU sets overarching legal obligations requiring platforms to identify, assess, and mitigate systemic risks arising from the operation of their services, including risks to electoral integrity. Platforms remain responsible for designing, implementing, and reporting on the measures used to address these risks.

Crucially, the DSA does not regulate individual pieces of content. It targets systemic risks arising from the design and operation of online platforms. Within the bounds of the law, individuals remain free to express false or mistaken beliefs. Risks to elections and democracy arise when platform features amplify such content at scale or when safeguards fail during sensitive electoral periods. Public debate often

¹ DSA Article 34.1(c) https://www.eu-digital-services-act.com/Digital_Services_Act_Article_34.html

² European Commission (2024). *Guidelines for Providers of Very Large Online Platforms and Very Large Online Search Engines on the Mitigation of Systemic Risks for Electoral Processes Pursuant to Article 35(3) of Regulation (EU) 2022/2065*. Brussels: European Commission.

³ An Coimisiún Toghcháin (2024). *Framework on Online Electoral Process Information, Political Advertising and Deceptive AI Content*. Dublin: An Coimisiún Toghcháin.

misrepresents this distinction, portraying EU regulation as censorship of speech rather than the mitigation of systemic risks. Nevertheless, it is through individual pieces of content and individual user activity that researchers, regulators, and the public can observe how platforms act, and how mitigation measures are implemented in practice. In addition, platforms' own terms of service govern the permitted content and behaviours for users of their services, with all major platforms outlining prohibitions on content and behaviour from their users that may affect elections and civic integrity.

The European Commission has set out recommended measures for platforms to address risks to electoral integrity. These include:

- signposting users to official election information,
- demonetising disinformation and reducing its visibility,
- enforcing rules against inauthentic accounts,
- providing tools to help users evaluate information,
- maintaining public repositories of political advertising, and
- preventing deceptive practices such as impersonation of candidates, manipulated media, fake engagement, non-transparent paid promotion, and coordinated inauthentic behaviour.

Some platforms participate in voluntary mechanisms to report their risk mitigation measures and to receive feedback from stakeholders. The EU Code of Conduct on Disinformation⁴ sets out commitments for platforms to address disinformation-related risks, including demonetisation, transparency of advertising, and cooperation with fact-checkers and researchers.

The Rapid Response System was first established by the Signatories of the Code of Conduct on Disinformation ahead of the 2024 European Parliament elections. It is “a time-bound dedicated framework of cooperation and communication” that allows organisations “to swiftly report time-sensitive content, accounts, or trends that they deem to present threats to the integrity of the electoral process and discuss them with the platforms in light of their respective policies”⁵.

During Ireland's 2025 presidential election, EDMO Ireland was invited to participate in this mechanism. Participating tech companies - Microsoft (LinkedIn), Meta (Instagram and Facebook), Alphabet/Google (YouTube) and TikTok - agreed to facilitate an expeditious response to any items flagged by EDMO Ireland and to offer

⁴ X (formerly Twitter) withdrew from the Code following its acquisition by Elon Musk and several companies have opted out of specific commitments. For example, Google (YouTube) and Microsoft (LinkedIn) have withdrawn from commitments to collaborate with fact-checking organisations aimed at ensuring users have access to verified information. See European Fact-Checking Standards Network (2025). [The Moment of Truth for the Code of Conduct on Disinformation](#).

⁵ European Digital Media Observatory (2024). [EDMO's assessment of the Rapid Response System of the Code of Practice on Disinformation](#). Florence: EDMO.

feedback on their decisions to act on or ignore the flagged items. The Rapid Response System has been utilised for subsequent elections, and is set to become a feature of the online content moderation process during elections in Europe. Against this background, this report examines the process and outcomes of the Rapid Response System during the Irish Presidential election of 2025.

The 2025 Irish Presidential Election

The Irish presidential election unfolded amid heightened concern about the safety of electoral candidates and politicians. In the preceding months, the Tánaiste and his family were subject to serious online and offline threats⁶ while Ireland's Digital Services Coordinator, Coimisiún na Meán, published its report on the online harassment of electoral candidates.⁷

Throughout 2025, Conor McGregor, a high-profile celebrity with a large Irish and international following, advanced a narrative that his failure to secure a presidential nomination would constitute a betrayal of democracy and the will of the Irish people. Following his withdrawal, elements of this narrative were taken up, to varying degrees, by other unsuccessful nomination campaigns.⁸ Elements of this framing were reflected in some aspects of a campaign encouraging voters to spoil their ballots. Spoiling a ballot is a legitimate form of political protest and a valid act of democratic participation. However, in this case, it was sometimes accompanied by false claims, including ones about the legitimacy of the nomination process and, by extension, the electoral system itself.

Early in the campaign, a coordinated smear campaign against one candidate was traced to an online actor, prompting legal correspondence with some platforms.⁹ In the final week of the campaign, an AI-generated video impersonated the leading candidate to falsely claim she was withdrawing from the race¹⁰. On the final day of the campaign, an online video depicting the alleged assault of female canvassers, which included the senior Irish politician. A man was subsequently arrested¹¹.

Taken together, the policy and regulatory context as well as unfolding campaign events set a clear expectation of vigilance and responsiveness from platforms.

⁶ O'Connor, N. (2025). *Gardaí launch investigation into 'significant' online threat to Tánaiste's family*. TheJournal.ie, 31 August.

⁷ Coimisiún na Meán (2025). *On the Digital Campaign Trail: Election Candidates' Online Experiences in the 2024 Elections*. Dublin: Coimisiún na Meán.

⁸ O'Brien, A. and Culloty, E. (2025). *From Free Speech to Violent Rhetoric: Elon Musk's X as a Threat to Electoral Integrity*. Dublin: EDMO Ireland.

⁹ Hickey, E. (2025). *Online smears spread about Jim Gavin came from 'far right' figure, Taoiseach says*. TheJournal.ie, 28 September.

¹⁰ McDonald, K. (2025). *AI-generated deepfake videos target presidential election*. RTÉ News, 22 October.

¹¹ Pepper, D. (2025). *Man arrested over alleged assault on two people canvassing with Mary Lou McDonald*. TheJournal.ie, 25 October.

Implementing the Rapid Response System in Ireland

By the time of the Irish Presidential Election in October 2025, the Rapid Response System had been activated for the European Parliament elections in June 2024 and for elections in France, Romania, Germany, Poland, Portugal, and Czech Republic. For the Irish Presidential Election, four companies¹² - Microsoft (LinkedIn), Meta (Instagram and Facebook), Alphabet/Google (YouTube) and TikTok agreed to facilitate an expeditious review of content flagged by EDMO Ireland and to offer feedback on decisions about that content. The Rapid Response System was activated one month before the election (29th September) and remained in place until a week after the election (4th November).

During this period, each company provided access to a reporting mechanism. Meta invited us to use its Official Complaints Portal, TikTok provided access to its Safety Reporting Tool Platform, and YouTube and LinkedIn invited us to report content directly to particular individuals. There were four online meetings, facilitated by the European Commission, between EDMO Ireland and platform representatives. These meetings provided an opportunity for the companies to clarify their decision-making process. It is important to note that EDMO could only report content that violated the platforms' own internal rules and terms of service, and the platforms were under no obligation to act on content flagged by EDMO Ireland.

Platforms' Terms of Service Versus the Digital Services Act

To ensure a rigorous and systematic approach to flagging content, EDMO Ireland first reviewed the terms of service as set out by each platform. In particular, we analysed rules regarding election integrity, disinformation, AI-generated content, political advertising, coordinated manipulation, harassment, violence, and monetisation/amplification of harmful content. These rules were assessed in light of stated commitments and expectations under the Digital Services Act. Overall, there appears to be a disconnect between the broad scope of the Digital Services Act and the narrow application of platforms' own rules.

For example, all platforms have rules about elections and civic processes. As summarised in Table 1, their terms of service describe the kinds of content that are not permitted and outline the actions that will be applied following breaches. Notably, the examples provided by the platforms focus almost exclusively on false claims about voting locations, voter eligibility, and participation interference. In contrast, the Digital Services Act has a more expansive understanding of what is

¹² As X withdrew from the EU Code of Conduct on Disinformation in 2023, it has not participated in the Rapid Response System.

required to protect public discourse and the fairness, transparency or legitimacy of electoral processes¹³.

Under the DSA and related frameworks, platforms' obligations regarding election and civic integrity concern: the ability of citizens to access accurate, reliable, and authoritative information necessary for participation in public debate¹⁴; preventing the dissemination and amplification of disinformation or misleading content that may misinform voter¹⁵; propaganda and falsely authoritative content, preventing coordinated manipulation, algorithmic amplification of harmful content, or misuse of platform systems¹⁶; maintaining risk mitigation measures are timely, proportionate, and effective throughout pre-electoral, electoral, and post-electoral periods¹⁷; preserving pluralistic public debate¹⁸, unfounded claims of electoral fraud, procedural flaws, interference, or institutional biases in favour or against certain persons, political parties or opinions, the large-scale dissemination of disinformation, misinformation and harmful conspiracy theories about candidates to public office, coordinated harassment campaigns (often targeting women and minorities), impersonation, and the circulation of synthetic media designed to mislead, and safeguarding against practices that undermine democratic participation¹⁹.

¹³ First report of the European Board for Digital Services in cooperation with the Commission pursuant to Article 35(2) DSA on the most prominent and recurrent systemic risks as well as mitigation measure (2025) <https://ec.europa.eu/newsroom/dae/redirection/document/121707>

¹⁴ DSA Article 34.1 (b) https://www.eu-digital-services-act.com/Digital_Services_Act_Article_34.html

¹⁵ DSA Article 35.1 https://www.eu-digital-services-act.com/Digital_Services_Act_Article_35.html

¹⁶ DSA Article 34.1(c) https://www.eu-digital-services-act.com/Digital_Services_Act_Article_34.html

¹⁷ DSA Recitals 83 and 84 https://www.eu-digital-services-act.com/Digital_Services_Act_Preamble_81_to_90.html

¹⁸ DSA Article 34.1(c) https://www.eu-digital-services-act.com/Digital_Services_Act_Article_34.html

¹⁹ *ibid*

Table 1: Overview of platform rules relating to elections

	Civic and Election Integrity Rules	Moderation Actions
YouTube	“Certain types of misleading or deceptive content with serious risk of egregious harm are not allowed on YouTube. This includes certain types of misinformation that can cause real-world harm, like certain types of technically manipulated content, and content interfering with democratic processes... Don’t post election-related content on YouTube if it fits any of the descriptions noted below. • Voter suppression • Candidate eligibility • Incitement to interfere with democratic processes • Keep in mind that this isn't a complete list.”	Content removal, account holder warning, three strikes over 90 days may lead to the termination of the channel. Severe abuse cases may lead to an immediate ban.
Meta	“In an effort to promote election and census integrity, we remove misinformation that is likely to directly contribute to a risk of interference with people’s ability to participate in those processes. We have additional policies intended to cover calls for violence, the promotion of illegal participation and calls for coordinated interference in elections”	Content removal, a strike may be applied against a user, utilising Meta’s “strike system” with repeat offenders receiving temporary restrictions.
TikTok	“We don't allow content that could mislead voters or interfere with elections, including false claims about how to vote, who can vote, or the outcome of an election”	“Flagged content that violates policy will be removed. Additional validation process for content posted by verified accounts. Accounts can be banned for repeated violations, a single severe violation, and avoiding enforcement”
LinkedIn	“Do not share false or misleading content: Do not share content that is false, misleading, or intended to deceive. Do not share content to interfere with or improperly influence an election or other civic process.”	“Depending on the severity of the violation, we may limit the visibility of certain content, label it, or remove it entirely. Repeated or egregious offences can result in account restriction”

Relatedly, all platforms have rules relating to AI-generated content (see Table 2). The rules primarily expect content creators to be responsible for labelling the AI-generated content they upload. In contrast, the Digital Services Act places clear obligations on the platforms rather than their users. In the context of elections, AI-generated content is recognised as a systemic risk²⁰.

More generally, the DSA asks platforms to apply “prominent marking” on generated or manipulated media that “appreciably resembles existing persons, objects, places”²¹. The Irish Electoral Commission also recommended that platforms provide tools for the prominent marking of election-related generative AI content²² and the DSA obliges platforms to tailor their election-related strategies in accordance with local authorities’ guidance²³.

²⁰ Guidelines for providers of VLOPs and VLOSEs on the mitigation of systemic risks for electoral processes <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52024XC03014&qid=1714466886277>

²¹ DSA Article 35(k) https://www.eu-digital-services-act.com/Digital_Services_Act_Article_35.html

²² AGS A Framework on Online Electoral Process Information, Political Advertising and Deceptive AI Content section B3 <https://www.electoralcommission.ie/media-release/an-coimisiun-toghchain-publishes-framework-on-online-electoral-process-information-political-advertising-and-deceptive-ai-content-in-advance-of-june-elections/>

²³ Guidelines for providers of VLOPs and VLOSEs on the mitigation of systemic risks for electoral processes Recital 3.4 <https://digital-strategy.ec.europa.eu/en/library/guidelines-providers-vlops-and-vloses-mitigation-systemic-risks-electoral-processes>

Table 2: Overview of platform rules relating to AI-generated content

	AI-Generated Content Rules	Moderation Actions
YouTube	“To help keep viewers informed about the content they're viewing, we require creators to disclose content that is meaningfully altered or synthetically generated when it seems realistic. Creators must disclose content that: • Makes a real person appear to say or do something they didn't do. • Alters footage of a real event or place. • Generates a realistic-looking scene that didn't actually occur.”	“When content is undisclosed, in some cases YouTube may take action to reduce the risk of harm to viewers by proactively applying a label that creators will not have the option to remove. Additionally, creators who consistently choose not to disclose this information may be subject to penalties from YouTube, including removal of content or suspension from the YouTube Partner Program.”
Meta	“Content digitally created or altered that may mislead. For content that does not otherwise violate the Community Standards, we may place an informative label on the face of content - or reject content submitted as an advertisement - when the content is a photorealistic image or video, or realistic-sounding audio, that was digitally created or altered and creates a particularly high risk of materially deceiving the public on a matter of public importance.”	“We may place an informative label on the face of content - or reject content submitted as an advertisement - when the content is a photorealistic image or video, or realistic-sounding audio, that was digitally created or altered and creates a particularly high risk of materially deceiving the public on a matter of public importance.”
TikTok	“Edited Media and AI-Generated Content (AIGC): We require clear labelling when AI or editing is used to realistically depict people or scenes. We don't allow AIGC that misleads about matters of public importance or that harms individuals.”	Unlabeled content may be removed, restricted, or labelled by our team, depending on the harm it could cause.
LinkedIn	“Do not share synthetic or manipulated media that depicts a person saying something they did not say or doing something they did not do without clearly disclosing the fake or altered nature of the material”.	“Depending on the severity of the violation, we may limit the visibility of certain content, label it, or remove it entirely. Repeated or egregious offences can result in account restriction”

Monitoring Election Content

EDMO Ireland established a monitoring procedure that would enable the identification of any election-related content potentially violating the platforms' own rules. As such, we were not simply flagging content that we deemed inappropriate, but content that we believe platform moderators should consider inappropriate by applying their own rules, commitments and standards.

Relevant content was identified through a collaborative, ethnographic approach to social media monitoring. Working with NewsWhip dashboards, EDMO Ireland monitored election content across multiple platforms and received reports from TheJournal FactCheck as well as specialist researchers from civil society organisations including the Hope and Courage Collective.

The process for monitoring election content typically involves identifying significant accounts, keywords, and hashtags across platforms. In most cases, it is limited to content that is publicly available. For example, the content shared within closed Facebook groups is inaccessible unless the monitor joins the group. It is also difficult to surface content from smaller communities or groups that are unknown to the monitors. Consequently, in the absence of comprehensive access to platform data, monitoring online content is an inevitably incomplete process.

EDMO Ireland triaged and evaluated all content against the platforms' terms of service and, when appropriate, flagged items via the Rapid Response System. Content that was particularly serious or time-sensitive was reported to the relevant platform by the organisation or individual that identified it, using the platform's standard reporting channels. This included in-platform reporting tools available to all users as well as direct reporting mechanisms used by journalists and civil society organisations where these were already established.

In total, 174 items were reported to the platforms as potential breaches of their stated rules and commitments to protect election integrity²⁴. Importantly, this figure only represents content that was observed and reviewed by the monitors. The total volume of content that could be considered in breach of platform rules is likely much higher.

The following pages provide a breakdown of the reported items.

²⁴ Originally, 175 items were reported. One item was removed during post-election data analysis as it was deemed outside of the election remit.

The Flagged Items and Responses

In total, 177 items were reported (see Table 3). LinkedIn had the lowest number of reported items (7 items) while Facebook had the highest (54 items). Considering the breakdown of the reported items, election-related disinformation was the most frequently reported type of content (85 items), followed by illegal content (49 items), AI-generated content (26 items) and harassment and abuse (13 items).

Disinformation items included false claims about the legality of the nomination process, false claims about the election process that were erroneously attributed to public officials, and assertions of systemic fraud in previous Irish elections. Illegal content primarily concerned images of ballots. In Ireland, it is illegal to share information about how you voted (including photos of ballots) until polling stations are closed²⁵. AI-generated content included a viral video impersonating one of the candidates as well as generic deepfakes. Harassment and abuse items included videos threatening and abusing political figures and canvassers.

Overall, platforms acted on half of the reported items. Actions taken included post removals, visibility restrictions, the addition of labels indicating AI-generated content, and account bans. Instagram was most likely to act on a reported item. It took action in 59 percent of cases. In contrast, LinkedIn did not act on any reported items.

Platforms' responses varied by item type. Platforms removed or geolocked all illegal content; that is, the images and videos of spoilt ballots posted on election day. In other cases, the platforms' application of their terms of services appeared inconsistent. Less than half of the AI-generated content reported by EDMO Ireland resulted in action such as the application of a content label. YouTube was the most responsive in this regard followed by TikTok. Facebook, Instagram, and LinkedIn did not act on any AI-generated content.

Only Instagram and TikTok took action on reported disinformation. Actions taken on election-related disinformation included removing content and account-level actions, as well as adding corrective labels to some AI content. Of the 13 items reported as harassment and abuse, platforms only acted on nine items. These items were reported to Facebook, Instagram, TikTok, and YouTube. TikTok acted on the three items reported to it while Instagram declined to act on the two items reported to it.

These actions are discussed in more detail in the following section.

²⁵Local Elections Regulations, 1995. (98.4) <https://www.irishstatutebook.ie/eli/1995/si/297/>

Table 3: Breakdown of reported items by type (items acted upon)²⁶

	# Items	LinkedIn	Instagram	Facebook	TikTok	YouTube
Election-related disinformation	85(13)	5(0)	11(1)	26(0)	30(12)	13(0)
AI-generated election disinformation	26(13)	2(0)	3(0)	2(0)	10(5)	9(8)
Harassment and abuse	13(9)	0	2(0)	3(2)	3(3)	5(4)
Incitement to violence	2(2)	0	1(1)	0	0	1(1)
Illegal content (e.g., spoiled votes)	49(49)	0	20(20)	21(21)	7(7)	1(1)
Political advertising	2(1)	0	0	2(1)	0	0(0)
Total	177(87)	7 (0)	37(22)	54(24)	50(27)	29(14)
% Actioned	49%	0%	59%	44%	54%	48%

In terms of the efficacy of the reporting mechanism as a rapid response, items were generally addressed within 24 hours. Despite having the lowest number of reported items, LinkedIn had the slowest response time, averaging about 20 hours. Many reported items remained active for days, which is surprising given the urgency associated with an election. TikTok had the fastest response time to flagged content (averaging 6 hours), with Meta coming second (averaging 8 hours), however, Meta declined to take action on the overwhelming majority of election-related content flagged during the Rapid Response System.

Before discussing the platform actions in more detail, it is worth making some observations about the participation of platforms in the Rapid Response System and their apparent preparedness for the election. During meetings, platform representatives often appeared unfamiliar with the context of the unfolding election and, as such, lacked context for discussing the reported items. Moreover, many were unable to articulate the rationale for content-moderation decisions. This is notable

²⁶ Note: In some cases, platforms applied multiple actions (e.g., removal and account-level restrictions).

given that there are evident inconsistencies in how those decisions were applied to similar pieces of content.

How platforms responded

Election-Related Disinformation

In total, 85 items were flagged as election-related disinformation. The platforms declined to take any action on the vast majority. Moreover, signposting to authoritative sources, corrective information, and other proactive information integrity efforts was poor across all platforms. For example, TikTok did provide an election fact-checking site, but it was not updated during the election. No corrective information was documented on Meta, YouTube, or LinkedIn in the run up to the election.

YouTube declined to act on any of the thirteen items reported to it. Nine of the reported items were posted on monetised YouTube accounts, meaning the creators were earning revenue from these videos. These videos included claims that the election was rigged and that candidates were barred from standing. YouTube stated that none of these posts violated its terms of service relating to election misinformation. This reflects a general pattern. Platforms typically declined to explain why a reported item failed to meet the threshold of a threat to election integrity. Instead, they simply noted that the item was considered to be in compliance with the terms of service.

Even when platforms did act, policy enforcement was often inconsistent. Five items were reported to TikTok for claiming that the “Irish Government banned candidates from running”. When reporting these items, EDMO Ireland provided TikTok with information from an authoritative source to explain why the claims were not true. TikTok removed one of the reported items, but not the other four. TikTok was unable to articulate why one post was deemed a violation of their terms of service, but not the others, despite the content being materially similar.




Figure 1: False claims casting doubt on the nomination process on 17/10/2025 Facebook.

In Meta's case, the inattention to corrective information is surprising given that it partners with factcheckers. In Ireland, TheJournal is Meta's fact-checking partner. Yet, Meta does not seem to have applied TheJournal's factchecks to obvious cases of disinformation. For example, a false story based on a fabricated quote attributed to a High Court judge originated on Facebook (see Figure 1). The quote called into question the constitutional validity of the nomination process. Six items with variants of this claim were reported to Meta along with a factcheck by TheJournal²⁷. Yet the platform repeatedly declined to take action against the posts, stating that it did not violate their terms of service, without elaborating further.

AI-Generated Election Content

Overall, 26 items were reported as AI-generated content that was potentially in breach of rules for failing to provide a label or for impersonating political candidates. Figure 2 is indicative of typical AI-generated content. Of the items flagged, 13 were actioned through applied labels, content removals, and, in one case, an account ban. There was no evidence of pre-emptive scanning and labelling of AI-generated content by the platforms. While some platforms applied labelling once content was reported, this was not consistent.

²⁷ The Journal (Oct 2025): Courts Service says quote about Áras nominations attributed to judge online was 'never spoken' <https://www.thejournal.ie/judge-party-whip-system-unconstitutional-presidential-election-6845937-Oct2025/>



Figure 2: An example of AI-generated content posted to both TikTok & YouTube on 07/10/25. After reporting YouTube added a “synthetic media label”, TikTok declined to take action.

Other platforms declined to apply any label or watermark when AI-generated content was reported. Meta noted that their platforms do not allow AI-generated impersonations of real people. However, four of the items flagged to Meta included AI-generated images of real people, but no action was taken on these posts.

TikTok’s enforcement of its policies regarding generative AI content was inconsistent. Of the ten pieces of AI-generated content reported to the platform, five were removed, while five received no action. Notably, a particular account was responsible for several AI-generated videos including depictions of Irish politicians (see Figure 3). While some videos were removed, others, such as the one in Figure 3, were not.




Figure 3: featuring AI-generated images of an Irish politician in handcuffs posted to TikTok on 06/09/2025. While TikTok removed several generative AI videos posted by this account, they declined to take action on 3 other clearly generative AI videos posted by the same account.

YouTube was the most proactive in responding to AI-generated content. It applied an “altered or synthetic content” label to seven of the nine reported items. However, it should be noted that YouTube’s label is not a watermark. If someone downloads the video - for example, with the intention to share it elsewhere - the label will not be included in the download.

In other cases, YouTube labels lack prominence in relation to the video. For example, Figure 4 depicts a screenshot from an AI-generated video that depicts election candidates and urges people to spoil their votes. While a corrective label has been applied (bottom left), it does not cover the visible image and is easily removed.




Figure 4: Poor placement of label warning about AI-generated content. This video posted to YouTube on 27/09/25. This video is one of a number that YouTube declined to remove despite featuring realistic portrayals of candidates.

YouTube did act swiftly to remove an account that shared highly realistic AI-generated videos targeting one of the candidates. Three days before the election, an AI-generated video purported to be a news report from the national broadcaster²⁸. The video impersonated one of the candidates announcing she was withdrawing from the campaign. The video was published by the account “RTENewsAI”, which appears to have been active on the platform since early September. The account was swiftly terminated after being reported, with YouTube referencing its policies prohibiting content that “makes a real person appear to say or do something they didn’t do or altering footage of a real event or place”. However, other content that also clearly violated this policy (see Figure 5) was clearly allowed to remain on the platform, albeit with labelling.

Before a piece of content (image/video/audio) was flagged by EDMO Ireland, it was assessed by an analyst trained in digital forensics and video/audio editing technology, with prior experience in large-scale investigations into Generative AI content on social media. The majority of the Generative AI content flagged by EDMO Ireland was of poor quality, with video and images often displaying telltale signs of

²⁸ RTÉ (Oct 2025) AI-generated deepfake videos target presidential election
<https://www.rte.ie/news/2025/1022/1539967-ai-generated-presidential-election/>

being generated by AI²⁹ (e.g. voice and audio artefacts, unnatural movement from subjects, morphs and jump cuts). Google Gemini’s AI detector³⁰ was used for additional verification in combination with Alliance4Europe’s AI4Trust platform³¹.

Harassment, Abuse, and Incitement to Violence

Election candidates and public representatives from across the political spectrum have been subjected to abuse and threats of violence. One incident occurring before the presidential election provides an illustrative example of how platforms address this issue of abuse towards public figures. In September 2025, a video posted on TikTok and X depicted a former senior Irish politician walking down a Dublin street. He is heckled by two men, who repeatedly call him a “traitor” and shout: “You shouldn’t be walking these streets ... Pushing lockdowns and vaccines—what about all the kids that were killed?” On X, the video was captioned: “Never let these traitorous animals walk the streets in peace... your days are numbered.” The incident was reported to the police and both men have now been charged with criminal harassment³². One of the people responsible was subsequently jailed for online posts targeting asylum seekers³³.




Figure 5: screenshot of a video harassing a former senior Irish politician Dublin Sept 2025

This video, along with excerpts from it, has been shared by multiple accounts across platforms. Many of these accounts routinely post about political and social issues

²⁹Van Ess (2025) A Reporter’s Guide to Detecting AI-Generated Content <https://gijn.org/resource/guide-detecting-ai-generated-content/>

³⁰ Kohl & Richardson (2025) How we’re bringing AI image verification to the Gemini app <https://blog.google/innovation-and-ai/products/ai-image-verification-gemini-app/>

³¹ AI4Trust <https://pilot.platform.ai4trust.eu/>

³² <https://www.rte.ie/news/courts/2026/0522/1574756-stalking-politician/>

³³ <https://www.irishtimes.com/crime-law/courts/2025/09/25/youtuber-becomes-first-person-jailed-over-posting-videos-of-asylum-seekers-online/>

and frequently target politicians with hostile or accusatory language, often labelling them as “traitors.” Despite clear evidence of harassment, the high public profile of the victim, the known history of the individual responsible, and media coverage of the incident, the video remained accessible on several platforms.

As indicated in Table 4, Meta and YouTube declined to act on our reports about these videos. Notably, the video was posted to Facebook by a verified account that was already at the centre of controversy surrounding “malicious smears” against a presidential candidate. The candidate, who subsequently dropped out of the race, issued legal letters to social media companies regarding false rumours about him and his family³⁴. A verified account means that Meta has confirmed its authenticity either because the account belongs to a notable person, brand, or entity, or because it is a paid subscriber to the Meta Verified service.

Table 4: Sample of accounts that shared an abusive video and platform actions after reporting

	Account Followers	Platform Action After Reporting
TikTok	N/A	Video & account removed
	12.6k	Video removed
Facebook	36k	No action
Instagram	26.7k	No action
	6.9k	No action
YouTube	8.2k	No action

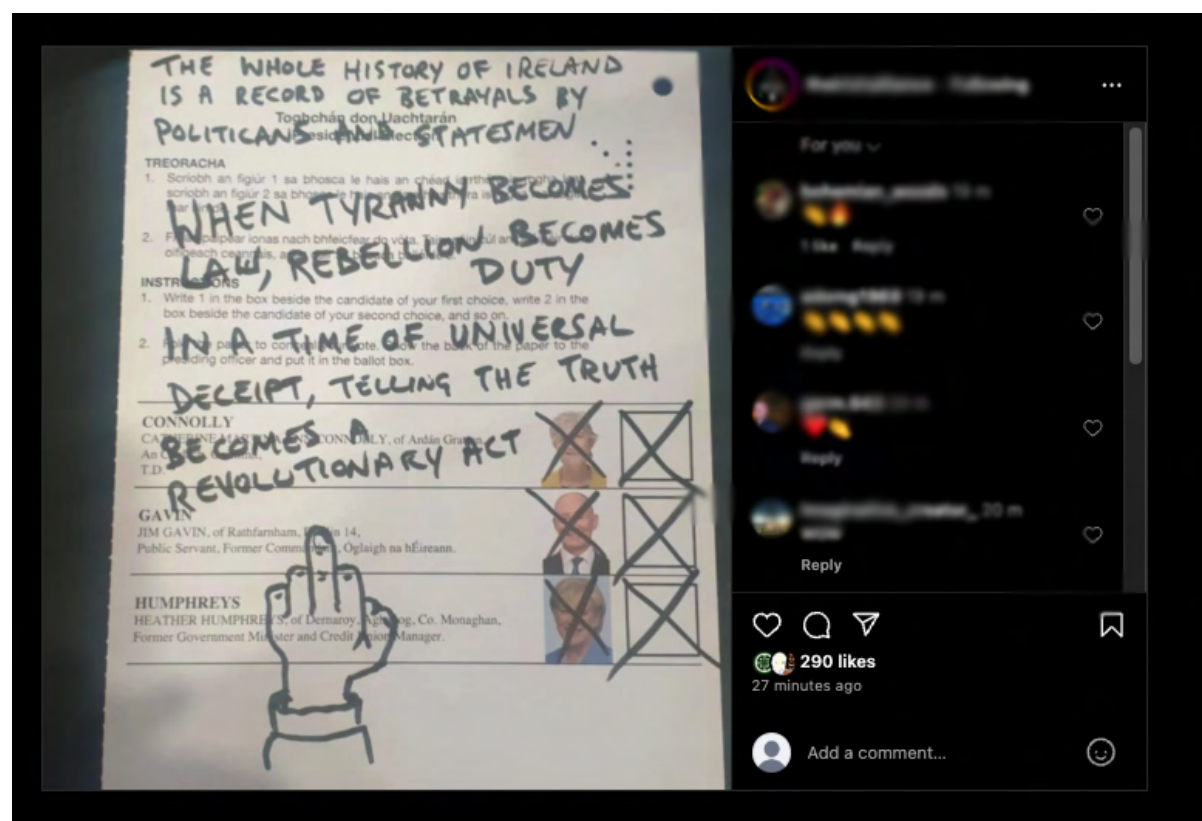
³⁴ <https://www.bbc.com/news/articles/c3vzn723l92o>

A second incident occurred on polling day, 24th October, as two volunteers and a senior Irish politician were canvassing in Dublin city. They were subjected to both a physical assault and sustained verbal harassment, and explicit threats of further violence, homophobic slurs and other abuse over the course of several minutes by a man who was subsequently arrested³⁵. However, this man was able to upload his video of the incident and share it on social media beforehand.

We reported multiple instances of the video to the platforms: TikTok removed the video, banned the account, and took further substantive action. YouTube and Meta, both of which declined to take action on the first video, swiftly removed six versions of this video published on election day.

Illegal Content

On polling day, Google, Meta, and TikTok all acted swiftly to remove or block the 49 items that were reported as illegal content; specifically, images of (spoiled) ballots. Under Irish election law, it is illegal to share information as to how you voted until polling stations are closed.



³⁵ <https://www.rte.ie/news/dublin/2025/1025/1540547-arrest-dublin/>

Figure 6: Sample spoiled ballot posted to Instagram on 24/10/25. 49 posts which include images or video of spoiled ballots were reported to platforms over the course of election day. Under Irish law, such images are illegal to distribute before polls are closed. All reports of illegal content were expeditiously removed.

While all reported instances of illegal content were removed, it is notable that there was no evidence of any effort to pre-empt illegal content. Platforms were not only forewarned of the potential issue of illegal content by EDMO Ireland ahead of polling day for the Irish Presidential election, but similar issues also occurred during recent the Moldovan³⁶ and Romanian³⁷ elections, with photos of ballots being disseminated on social media on polling days for their respective elections. Not only did none of the platforms have policies in place for this eventuality, only one platform articulated its policy position on content relating to the “spoil the vote” campaign. Unfortunately, this was provided on 29th October, five days after the election.

Every report of illegal content made to platforms was actioned, with the content uniformly and swiftly removed or geoblocked by every platform. This is at odds with the, less than systemic, application of the platforms' policies regarding other reported content that was often in clear violation of their own terms of services.

Account-Level Actions

Flagged items were consistently assessed in isolation from the account that created or shared it, including the prior conduct of the account holder. For example, nine of the items reported to YouTube were published by monetised accounts. As noted, items flagged to Facebook came from accounts that had previously been reported for illegal and harmful behaviour. By viewing each item in isolation from the broader pattern of accounts, the nature of systematic risks is lost. Decisions appeared to focus narrowly on individual items of content, rather than on how platform systems might contribute to the visibility, spread, or financial incentivisation of disinformation.

TikTok was the platform most likely to ban accounts. TikTok allows users to “create multiple accounts”, but it “may restrict or permanently ban” users that “create a new account after your account has been banned for a severe violation”, “post violative content on a new or existing account”, or “spread content violations across multiple accounts”. In practice, it seems account bans are easily circumvented. Users can recreate accounts almost immediately, often under slightly altered

³⁶ <https://hotnews.ro/video-un-numar-neobisnuit-de-mare-de-alegatori-si-au-fotografiat-buletinul-de-vot-dupa-ce-au-iesit-din-cabina-topul-neregulilor-de-la-alegerile-din-republica-moldova-de-ce-au-fos-2074906>

³⁷ <https://www.europafm.ro/politia-romana-face-verificari-dupa-ce-romani-din-diaspora-si-au-filmat-votul-si-au-publicat-imaginile-pe-retelele-de-socializare/>

usernames or through new device registrations. The failure to detect and address repeated re-registrations or coordinated use of multiple accounts suggests weaknesses in enforcement mechanisms. Such evasion enables the continued spread of abusive, misleading, or manipulative content, undermining TikTok's stated policies and raising questions about the platform's effectiveness in mitigating systemic risks during election periods. Ultimately, TikTok did acknowledge the need for measures to prevent the circumvention of account bans. However, at the time of writing, six months later, two of these individuals appear to be active on the platform circumventing TikTok's ban.

Political Advertising

The EU adopted the Regulation on the Transparency and Targeting of Political Advertising, designed to prevent manipulation and interference in elections. The regulation requires platforms to label political ads, restrict targeting, and maintain searchable archives for several years. Rather than comply with these obligations, Google and Meta both decided to stop accepting political adverts in the EU: Google on 22 September and Meta on 1 October. Despite this, enforcement has been inconsistent; Google has reportedly accepted political ads in Hungary despite its declared ban.

There is a lack of access to political advertisements surrounding the 2025 Irish Presidential Election and the period leading up to it. Google has deleted the archive of political ads and Meta has removed records of at least 83 campaign ads. This gap represents a significant setback for transparency and accountability in digital campaigning. Two queries relating to political advertising were raised with Meta. One related to the deletion of one candidate's ad library after they withdrew from the election³⁸. This was attributed to a technical error and later restored. EDMO Ireland felt it necessary to express their frustration to Meta at the lack of data access and reporting mechanism available during the Rapid Response System to monitor and report issues relating to their ad library and advertising policies. A second more general query regarding enforcement of Meta's ad library was not answered, although it was submitted at a late stage in the Rapid Response System process.

Conclusion

The establishment of the Rapid Response System represents a positive and necessary step forward. As a dedicated mechanism for flagging potentially harmful content

³⁸ Irish Times (Oct 2025) Deletion of Jim Gavin ads shows we can't rely on big tech for democratic architecture <https://www.irishtimes.com/opinion/2025/10/07/political-ads-deletion-shows-we-cant-rely-on-big-tech-for-democratic-architecture/>

during time-sensitive election periods, it provides platforms with a structured channel to receive and act upon urgent referrals. The existence of this system is a valuable foundation, demonstrating a shared commitment to electoral integrity and timely intervention.

That said, the process remains nascent and requires further development. A foundational issue lies in the persistent disconnect between platforms' terms of service—which serve as their primary guidebook for content moderation decisions—and the broader systemic risks that EU regulatory frameworks seek to address. While the Rapid Response System facilitates individual content flagging, it does not yet bridge the gap between case-level enforcement and the systemic accountability envisaged by recent legislation.

Furthermore, despite the insights generated through the process, little is known about how platforms actually reach their moderation decisions. This opacity leaves external observers in the dark regarding the consistency, rationale, and possible divergences in platform responses. Without greater transparency, it is impossible to assess whether decisions are applied uniformly or whether inconsistencies undermine the system's effectiveness.

In sum, the Rapid Response System is a welcome and useful instrument, but it cannot substitute for genuine accountability. Platforms need to do much more to meet their obligations to protect the integrity of elections.