

Understanding online information-based threats to elections

13 October 2026

Background

This report is one of a series of advice notes published by Ofcom’s Online Information Advisory Committee (‘the Committee’), a body established under Section 152 of the Online Safety Act. The Committee’s function is to provide independent advice to Ofcom on how providers of regulated services should deal with disinformation and misinformation on their services, and on Ofcom’s exercise of certain powers and duties as they relate to disinformation and misinformation on regulated services.¹

The Committee’s approach is grounded in respect for fundamental rights, particularly the right to freedom of expression under Article 10 of the Human Rights Act.

The aim of the Online Safety Act regime is to reduce harms suffered by people in the UK when they are online, while simultaneously protecting freedom of expression. The Committee’s work reflects this scope, and any advice should be read as relating to UK users. The Committee’s work covers areas where the harms relate to the production and distribution of false and misleading information, as outlined under the UK’s Online Safety Act. This is a diverse field of potential harms, only some of which are illegal. The Committee’s work examines a series of specific areas of harm and provides advice in relation to them.

The first advice note focused on the intersection of disinformation, misinformation and online fraud.² This second report focuses on information-based threats to the integrity of elections: the malicious use of false, misleading, or deceptive information, and deceptive behaviour, to influence electoral processes, participants, or public confidence in elections. The public reports substantial exposure to false or misleading electoral content. Ofcom’s General Election tracker in 2024 found that a majority (60%) of respondents said they had encountered information about the General Election that they felt might have been false or misleading at least once in the prior week.³ The Electoral Commission’s report on the 2024 UK General Election found that many respondents said they saw misleading or inaccurate information about political parties’ policies (61%) and candidates (52%), while some (30%) saw misleading or inaccurate information about how elections work, and around a quarter (24%) reported seeing or hearing a deepfake about the election.⁴ In the 2025 UK local elections, misinformation and disinformation were the third largest worry among survey respondents, even as public confidence in the elections remained high.⁵ We have also heard warnings about the risks of misinformation and disinformation in the electoral context from various Parliamentary bodies.⁶

¹ This includes Ofcom’s transparency powers under the Online Safety Act, s.77, and Ofcom’s media literacy functions under the Communications Act 2003, s.11.

² Online Information Advisory Committee, 2025. [Understanding Online Financial Harm](#).

³ Ofcom, 2024. [UK General Election news and opinion-formation survey 2024](#), p. 9.

⁴ Electoral Commission, 2024. [Report on the 2024 UK Parliamentary general election and the May 2024 elections](#).

⁵ Electoral Commission, 2025. [Report on the May 2025 local elections in England](#).

⁶ See, for example, House of Lords Library (Thomas Brown), 2026. [Threats to UK democracy: Disinformation, foreign interference and declining public trust](#).

Given these concerns, information-based threats to elections are an appropriate topic for this Committee's advice.

The Committee asked Ofcom to collate relevant literature on the tactics, techniques and procedures associated with online information threats to elections, and on the quality of the evidence for each. Ofcom also brought together information about countries whose recent electoral experience is instructive for the UK. The Committee is publishing the literature review associated with this report so that readers can better understand the underlying research that informs our advice.⁷

The members of Ofcom's Online Information Advisory Committee wrote this report independently, drawing on the literature review and their own relevant expertise.

Several points about scope should be made at the outset. First, offline information threats in the UK are not covered here. Dangerous falsehoods circulated in a pub, while conceivably capable of leading to harm, do not fall within scope.

Second, our central concern is with *disinformation*. The UK government has defined disinformation as the deliberate creation and spreading of false and/or manipulated information that is intended to deceive and mislead people, either for the purposes of causing harm, or for political, personal or financial gain. Misinformation is defined as the inadvertent spread of false information.⁸ The line between these categories is not always clear; what begins as malicious disinformation may soon become innocent misinformation. Further, regulated services perennially struggle to distinguish them due to the difficulty of inferring a speaker's mental state (a difficulty faced by both human and machine moderators). But we seek to set aside the latter category. Political expression is core to freedom of speech under Article 10. Robust debate will inevitably involve inaccuracies, exaggerations, simplifications, and honest mistakes of fact; such errors are perfectly ordinary parts of public discourse, protected by the freedom of speech. There is not, and cannot be, a general prohibition on falsehood in an open society. Our concern here is with the malicious, intentional use of falsehoods to damage electoral institutions. This includes falsehoods about a speaker's own identity; cases in which a speaker inauthentically misrepresents their own identity fall within our scope, even if the speaker isn't promulgating other falsehoods.

Third, some of the conduct canvassed in this note overlaps with other categories of harmful (and often illegal) content. Harassment and intimidation of candidates, non-consensual intimate image abuse, foreign interference, cyber-attacks on electoral infrastructure, and fraud are all addressed by other parts of the criminal law and (only some within the Online Safety Act regime,) and in some cases by other regulators. This note considers such conduct only insofar as it operates through false, misleading, or inauthentically presented information.

⁷ Ofcom, 2026. [A literature review of online information-based threats to elections](#).

⁸ This broadly tracks the usage of the terms by the UK government. See, for example, Cabinet Office and Department for Science, Innovation and Technology, 2023. [Fact Sheet on the CDU and RRU](#). As quoted in Communications and Digital Committee, 2024. [The Future of News](#), chapter 7.

Finally, it is imperative to stress the limits of the available evidence in this area. The literature on information threats to elections is large, fast-moving, and highly uneven. For many widely discussed threats there is little or no high-quality evidence. For almost none is there robust evidence about effects on voting or other behaviour. What’s more, a substantial body of research argues that the prevalence and impact of electoral disinformation are overstated. Accordingly, confident public claims about harm can themselves do damage. Such claims pose risks to freedom of expression, by inviting disproportionate responses, and to confidence in elections, by amplifying narratives of electoral mischief that malicious actors themselves seek to spread.⁹ The Committee has sought to reflect this uncertainty honestly throughout, and its advice to Ofcom is shaped by it.

The next section of the report critically reviews how these harms manifest online, drawing on the literature review. The subsequent section assesses the existing legal, regulatory, and media literacy tools for confronting these harms, clarifying which belong to Ofcom and which do not. The final section presents the Committee’s conclusions and recommendations to Ofcom on how to confront these challenges while upholding freedom of speech.

How harms manifest online

The literature review usefully organises online information-based threats to elections into a taxonomy of four overlapping categories: information manipulation, ecosystem manipulation, marketing manipulation, and participant manipulation. The taxonomy describes what has been attempted and observed; as discussed later in this section, it should not be read as a hierarchy of impact, which the evidence cannot support. Table 1 summarises the taxonomy.¹⁰

Table 1: Literature-derived categorisation of information threats to elections

Category	Aim	Illustrative techniques
Information manipulation	To influence citizens’ beliefs, preferences, perceptions, narratives, or understanding	Fabricated or altered content (including fabricated news outlets, deepfakes, fabricated or misrepresented statistics and polls); deceptive forms of narrative shaping (conspiracy theories, certain polarising language or rage bait); delegitimisation of electoral processes.
Ecosystem manipulation	To influence the discoverability, delivery, or apparent provenance of information	Search and algorithmic manipulation; exploitation of data voids; information (over)saturation; LLM ‘poisoning’; bots, fake accounts and coordinated inauthentic behaviour; astroturfing; inauthentic websites, proxies and information laundering; synthetic personas; technical disruption (e.g., hacking and ransomware to disrupt or confuse).

⁹ See the Methodology and Limitations sections of the literature review for the search strategy, inclusion criteria, and a full account of the constraints on the evidence base.

¹⁰ Because our aim is to offer a high-level synopsis that highlights some striking findings in the literature, we do not cite the full details here of the studies referenced in the accompanying literature. The full citations, of course, can be found there.

Marketing manipulation	To influence electoral outcomes through illicit or opaque funding or the manipulation of advertising	Microtargeting of tailored narratives; paid or covert use of influencers; fake endorsements; inauthentic advertising; “dark PR”; circumvention of political advertising restrictions and spending limits
Participant manipulation	To influence participation in elections rather than beliefs or preferences	Voter suppression content; automated persuasion; cognitive warfare aimed at inducing disengagement; harassment and trolling of voters; gendered disinformation targeting candidates; deepfake impersonation of candidates; non-consensual intimate image abuse; doxxing; strategic lawsuits against public participation

Information manipulation

The first category identified in the literature review is information manipulation, which aims to shape voters’ beliefs, preferences, perceptions, narratives, or understanding through the content of what is communicated. The literature review documents several sources reporting on examples of information being fabricated or altered in an attempt to influence elections. One study found false and decontextualised content to be the predominant form of electoral disinformation in the 2024 European Parliament elections. In the 2024 UK General Election, one source found that a disinformation network generated inauthentic news sources, combining fake articles with real news stories, re-written to include ‘pro-Kremlin’ narratives; another found that inflammatory fake content was disseminated in Scotland in Gaelic, presumably to stoke discontent between speakers and non-speakers, and that such content is likely to circulate undetected for longer periods of time due to being spoken by a relatively small population. Misleading claims by mainstream political actors are also common across the political spectrum.

Deepfakes have attracted particular attention. In the final days of the 2023 Slovak election campaign, a fake audio clip appeared to show the politician leading the polls discussing electoral fraud; his rival won, though a study describing the episode considered it reductive to attribute the result to a deepfake alone. In the UK, Full Fact concluded from the 2026 local elections that AI imagery is now a real factor in UK elections, observing fake videos of voters, fabricated meetings and rallies, and altered images of real events. Yet scale matters: one study identified just 16 confirmed viral cases of deepfake disinformation during the 2024 UK general election, and 11 across the 2024 EU and French elections combined—and it is unclear whether deepfakes meaningfully affected results.

Related phenomena complicate the picture further. Sometimes, campaigns produce so-called ‘soft fakes’ of their own candidates in ways that blur the distinction between manipulation and self-promotion. There is increasing attention to the so-called ‘liar’s dividend’, by which genuine content gets dismissed as fabricated. One experimental study in the US found that casting doubt on negative accusations is more effective in maintaining support for a candidate than either apologising or staying quiet.

AI-powered large language models (LLMs) are an increasing focus of research. One study found that almost all of the LLMs assessed could be used to produce authentic-seeming and often localised election disinformation, including about voter ID requirements, polling arrangements, and MP

behaviour. At the same time, observation of recent elections found minimal AI-generated false content in the 2024 European Parliament elections, and only one identified case of deceptive campaign adverts made by generative-AI in the 2024 UK general election. The upshot, then, is that the capability of AI-generated disinformation plainly exists; but widespread consequential use is not evidenced, at least not yet.

Within this category, attempts to create the delegitimisation of electoral processes stands out as the best-evidenced harm. In a study of the 2024 European Parliament elections, electoral integrity was the most prevalent topic of debunked falsehood—over 20% overall and over 80% in Germany. In the run-up to the 2026 Bulgarian parliamentary elections, coordinated attempts cast pre-emptive doubt on the legitimacy of the electoral results; according to one study, three information operations gained over 340 million impressions across coordinated media and social networks. Rather than convince voters of any particular views, these efforts were reported to have served to erode trust before voting day. Related survey research finds that exposure to claims of voter fraud reduces confidence in electoral integrity, and that losing voters are more likely to view the process as unfair and become more conspiratorial after elections. One study, drawing on more than 80,000 respondents across 74 democracies, stresses the association between disinformation campaigns and less accurate, more polarised beliefs about election fairness.

Notably, one large US study found that official messaging about trust in elections can be strongly effective at increasing trust; accordingly, the absence of visible, authoritative official information may itself be a vulnerability.

Foreign interference cuts across this category. The Council of Europe reporting describes a sharp increase in the scale of foreign information manipulation and interference (FIMI) in European electoral processes. One study of Russian FIMI activity found elections to be its primary target and described a playbook of delegitimising political leadership before elections, seeking to exacerbate domestic divisions during campaigns, and sowing doubts about electoral fairness. In the UK, the Rycroft Review has considered foreign interference in British politics to be both “real and persistent.” The Prime Minister’s recent announcement of a new National Centre for Information Defence, about which we are keen to learn more, reflects this concern.

Ecosystem manipulation

The second category identified in the literature review is ecosystem manipulation, which aims to influence the discoverability, delivery, or apparent source of information. The literature documents search engine and algorithmic manipulation for electorally relevant queries—seeking to boost preferred content so that more people see it, more frequently. It also documents the exploitation of ‘data voids’, where a lack of existing organic user-generated content around certain terms provides an opportunity to fill the void. Fake accounts and bots can also engage in information saturation, saturating a service with new content to drown out or obscure disfavoured content. One survey showed that these strategies can induce fatigue, cynicism, and disengagement from democratic processes.

The coordinated use of inauthentic accounts and information is among the most consistently documented threats in the literature. One 2026 study discovered over 30,000 inauthentic accounts across seven major online platforms. In the 2025 German federal election, a coordinated network of 50 accounts, amplified by a bot network of over 6,000 accounts, spread false claims in realistic videos

impersonating news outlets, law enforcement, and academics. Provenance-obscuring infrastructure is also well documented: the Doppelgänger network has been reported to imitate news organisations while promoting Russian narratives; the Storm 1516 network allegedly sought to influence Germany's 2025 elections with more than 100 fabricated websites spreading falsehoods, and it reportedly also ran a coordinated influence operation during the 2026 Hungarian election. In numerous cases, such behaviour operates through domestic proxies, such as pseudo-thinktanks, and information laundering schemes – such as tricking legitimate news media to report disinformation as fact.

While traditional forms of 'coordinated inauthentic behaviour' involve human-staffed troll farms, new AI technologies are making it possible to run these coordinated efforts through bot networks. Synthetic personas make these efforts even more potent. In the run up to the 2026 Hungarian election, a network of nearly 100 AI-generated personas spread content supporting the governing party.

LLM chatbots themselves constitute a further potential target of manipulation through so-called poison attacks. One worrying trend is 'LLM grooming' — flooding the internet with content potentially intended to be ingested by models so that they echo favoured narratives. This trend is exemplified by the pro-Kremlin Pravda network, much of whose outputs appear designed for LLM training and search engine crawlers, rather than human readers. But the evidence here is mixed: one study found popular LLMs repeating Russian disinformation about Ukraine, while another assessment found no consistent recapitulation of Russian propaganda. Thus, at present, based on evidence examined, there is no conclusive evidence on the effectiveness of LLM grooming as a genuine information threat to elections.

A final way to manipulate the information ecosystem is through direct technical disruption — such as the surge in ransomware attacks against Moldovan government websites and news portals during the 2025 elections. These cybersecurity threats can be coupled with physical sabotage and kinetic attacks. This mechanism sits at the edge of this Committee's remit. It is nevertheless relevant, because even minimal cyber-attacks can have a disproportionate psychological effect once publicised, a process called 'perception hacking'.

Marketing manipulation

The third category identified in the literature review is marketing manipulation, which aims to influence elections through illicit or opaque funding or the manipulation of advertising, including otherwise legitimate techniques conducted without sufficient disclosure. The literature documents microtargeting of tailored narratives at specific audience segments. Influencers are sometimes (covertly) paid to disseminate political messaging: such techniques were reported in the annulled 2024 Romanian presidential election, where around 100 influencers amplified emotionally charged and nationalist narratives coordinated and financed through opaque channels, and in Hungary, where the ruling party was reported to have spread disinformation through a coordinated pro-government influencer network. Endorsements can also simply be faked, as when the Irish Tánaiste was impersonated in a deepfake endorsement. Inauthentic advertising is cheap at scale: one experiment generated over 200,000 views and more than 17,000 inauthentic comments for a total cost of €252.

The occasional bans on political advertising adopted by some platforms were reported in some literature as not always well enforced. After one platform's ban, 457 political ads, including deepfakes, were identified in the first six weeks of the 2026 Hungarian campaign. Further, actors may seek to circumvent campaign spending limits or advertise on networks that do not enforce disclosure requirements.

Despite the activity in this space, the effectiveness of marketing manipulation during elections is plainly inconclusive. This is, in part, because the effectiveness of online political advertising as such is disputed. One US study during the 2020 presidential election of over 60,000 users on Facebook and Instagram tested for the effects of removing political ads on political knowledge, polarisation, perceived electoral integrity, political participation, candidate favourability, and turnout. It found no detectable effects.

(To foreshadow our later discussion: tools for dealing with this third category – marketing manipulation – fall primarily to bodies other than Ofcom. Campaign finance and digital imprints are matters for the Electoral Commission; meanwhile, paid-for advertising is the remit of the Advertising Standards Authority, and it excludes *political* advertising from its remit.)

Participant manipulation

The final category identified in the literature review is participant manipulation, which aims to influence participation in elections rather than beliefs or preferences (though this overlaps with that first category, since persuasion is an obvious way to influence electoral behaviour). On the voter side, one large US study of exposure to targeted advertising identified various types of voter suppression content — including election boycott messaging (“don’t vote!”), and deception about voting logistics. The study found that, on average, those exposed to voter suppression ads were less likely to turn out to vote, an estimate the study extrapolated to millions of votes nationally—though this remains a single study and should be treated with corresponding caution.

On the candidate side, there is evidence that candidates receive various forms of intimidation and harassment. In the 2024 UK general election, the Electoral Commission reported that many candidates, especially women and ethnic minorities, experienced intimidation and harassment, and over half of candidates responding to its survey said that fears of abuse motivated them, at times, to avoid campaigning. In the 2024 Irish General Election, 55 incidents of politically motivated violence, threats, harassment and smears were observed in the final five weeks of campaigning. Women are particularly vulnerable. In an analysis of comments under TikTok videos posted by candidates for the 2024 European Parliament elections, women candidates received substantially more hateful, defamatory, derogatory or discriminatory comments than their male counterparts. Non-consensual intimate image deepfakes disproportionately target women politicians. Deepfakes have also been deployed to sow confusion; in Ireland, a fake news report falsely alleged that a top candidate was withdrawing from the race.

Crucially, whether these attacks on voters and candidates aggregate to alter the outcome of elections is not remotely substantiated by the evidence. Even so, the harms to those individuals affected—voters who show up to the polls on the wrong day, or candidates intimidated into silence—are real, and are cause for concern even if their aggregate effects remain unknown.

The limits of the evidence

In reviewing the findings of the literature review, we have endeavoured to recognise the real harm that can be caused by maliciously deceptive conduct online, while also recognising that the known impact of such behaviour can be limited. It is worth re-emphasising those limitations. According to one review, estimates of exposure to falsehoods as a share of online news and civic content hover around 1–10%, albeit rising for contentious topics. Other research argues that average exposure to such content is low, and that it is concentrated among a narrow group motivated to seek such content out; that people mostly consume falsehoods they are already disposed to believe; and that engagement should not be conflated with belief, nor belief with behavioural change. The widely predicted wave of AI-driven electoral manipulation in the elections of 2024 largely did not materialise, with one review stating that it was largely overshadowed by more familiar factors — including politicians’ own willingness to mislead.

Combating electoral disinformation

There is a variety of legal and regulatory regimes relevant to the harms discussed in this report and this is not a comprehensive summary. But the starting point is different from the Committee’s first advice note, which focused on financial disinformation in the form of fraud. Fraud is a clear and uncontroversial category of criminal content. Fraud offences are priority offences under the Online Safety Act, so fraudulent content is priority illegal content. In contrast, it is not necessarily a crime to spread falsehoods in the context of an election. As we stressed earlier, inaccuracies and honest mistakes are a normal part of political debate, and are not illegal in the UK. And while it is clear that the malicious, intentional spreading of falsehoods to influence or undermine elections is morally objectionable, there is no single overarching criminal offence that captures all such disinformation. The tools for combating electoral disinformation, then, are multiple, targeted at specific subcategories, and distributed across many UK institutions. Here we set out the landscape in detail, clarifying whose role it is to do what. This will, in turn, clarify what Ofcom can, and cannot, reasonably be expected to do in response to this challenge.

First, consider the various criminal offences bearing on electoral disinformation and related harms. These include the Representation of the People Act 1983, which prohibits false statements of fact about a candidate’s personal character or conduct for the purpose of affecting an election;¹¹ and the National Security Act 2023, which prohibits interference in elections on behalf of foreign powers.¹² There are also criminal statutes prohibiting harassment, threats, stalking, and the sharing of intimate images without consent—none of which are election-specific, but are clearly relevant for preventing the kinds of candidate intimidation we have discussed. The new false communications offence, which prohibits sending information known to be false with an intention to cause non-trivial psychological

¹¹ [Representation of the People Act 1983, s.106](#). Notably, this does not extend to false statements about a candidate’s politics. [The Elections Act 2022](#) restated and clarified the offence of undue influence.

¹² [National Security Act 2023, s.13-15](#).

or physical harm, is also potentially relevant.¹³ And the hacking of electoral infrastructure violates the Computer Misuse Act 1990.¹⁴

These offences are investigated by the police and prosecuted by the Crown Prosecution Service—not by Ofcom. Still, Ofcom has a vital role to play under the Online Safety Act with regard to specified illegal content. Under the Act, regulated services are required to use proportionate systems and processes to protect users from illegal content, with enhanced proactive duties in respect of designated “priority offences”—including foreign interference. Ofcom’s role is to ensure that regulated services comply with these duties. Regulated services must also enable users to report suspected illegal content, and must swiftly take down illegal content once they become aware of it.

We thus advise Ofcom to ensure that regulated services are complying with their duties to take action against illegal content, especially priority illegal content.

Second, the Online Safety Act empowers Ofcom to enforce additional duties on categorised services. These include duties of transparency reporting (which apply to all categorised services) and, for Category 1 services, duties to enforce one’s own terms of service consistently. These duties are relevant precisely because many Category 1 services have themselves voluntarily adopted terms of service against electoral misinformation, fake or impersonating accounts, coordinated inauthentic behaviour, and deceptively manipulated media. These terms sweep much more broadly than the criminal law. Under the Act, services are expected to live up to the commitments they make to their users on what is and isn’t allowed; and one of Ofcom’s roles is precisely to hold them accountable. One way for Ofcom to limit the harms of electoral disinformation, then, is simply to ensure that services are abiding by their own terms of service. In fact, given that many services are complying with their illegal content duties by adopting broad terms of service that encompass illegal content, holding platforms to their terms of service is all the more important. Note that Ofcom’s additional powers involve holding Category 1 services accountable for protecting journalistic content and content of democratic importance – and for having particular regard to the importance of protecting users’ right to freedom of expression within the law which can militate against overenforcement, ensure consistency and transparency, and help to protect free speech online.

We thus advise Ofcom to maintain its schedule to finalise Codes of Practice and guidance regarding Category 1 services’ additional duties. This will clarify how services will be held accountable for enforcing their own terms of service, including on mis/disinformation should they choose to have policies for them.

Third, Ofcom has statutory media literacy functions under Section 11 of the Communications Act 2003. These functions are particularly significant in the context of elections. As we have seen, the vast preponderance of electoral misinformation is entirely legal.

We advise Ofcom to continue its media literacy work in educating UK users on how to spot unreliable sources and falsehoods in the context of elections, especially by leveraging its supervisory relationships to urge platforms to implement recommendations from its voluntary guidance on best practices.

¹³ [Online Safety Act 2023, s.179.](#)

¹⁴ [Computer Misuse Act 1990.](#)

Fourth, Ofcom can commission new research to help fill some of the gaps in existing knowledge. Claims of harm in this area frequently run ahead of the evidence. This is a problem for many reasons. It can lead us to overreact, embracing unnecessary or disproportionate remedies that undermine freedom of expression. Further, exaggerated public claims about a corrupted public sphere do the delegitimisers' work for them, feeding both the liar's dividend and the perception hacking phenomena discussed above. New research, then, is needed. The literature review identifies limited evidence on the effects of information threats on voting and other behaviour; limited longitudinal research; and limited research on private messaging services, which may be trusted differently by users and may be used as a laundering point by malicious actors.

What's more, current research overwhelmingly studies the environment we have had: adult users, on large public social media services, in past elections. But the environment ahead may look quite different. In the UK, we will likely soon see the extension of the franchise to sixteen- and seventeen-year-olds in reserved elections under the Representation of the People Bill; minimum-age restrictions on social media have been introduced in some jurisdictions and are actively debated in others; and generative AI tools are being rapidly adopted as sources of information and opinion. It is possible to foresee a cohort of first-time voters who are restricted from the very platforms researchers know how to study, and who instead learn about politics substantially through AI intermediaries whose electoral reliability is questionable and whose influence on political belief is almost entirely unresearched. Some of the most important challenges for electoral information integrity, in other words, are ones for which there is very limited research — and the research base we have may be least relevant to the voters who come next.

We advise Ofcom to commission new research and to work with UKRI to fill the gaps in our current knowledge base. Particular gaps include (1) the most effective media literacy strategies for people under eighteen; (2) understanding how people are using generative AI to seek out political information, the ways in which mis/disinformation can emerge in generative AI, and how generative AI use is influencing voting beliefs; and (3) how to foster trust in authoritative information during an election.

Notwithstanding these recommendations, it is vital to stress the limits of what Ofcom can achieve. Arguably, the pressure on Ofcom to act is at its peak during elections. Yet it does not have any special powers to intervene in electioneering or other discourse during an election. Further, Ofcom has no role in making decisions about individual posts or accounts; it cannot require specific pieces of content to be taken down. Beyond cases of fraud, paid-for advertising is broadly outside the scope of the Act. And the impartiality and accuracy rules that Ofcom enforces for broadcasters don't extend to online service providers under the Online Safety Act. Finally, it bears repeating that the vast preponderance of political falsehood is protected speech. The default remedy when encountering false content online is, inevitably, to engage in counter-speech in an effort to debunk it, or simply to ignore it.

Importantly, Ofcom is not remotely the only actor in the landscape responsible for combating the varieties of information-based threats to elections.

First, the Electoral Commission regulates campaign finance and spending. It also administers the digital imprints regime introduced by the Elections Act 2022, which requires digital campaign

material to declare who is promoting it.¹⁵ And it provides authoritative public information about electoral processes—something that research suggests is important, since official messaging about elections can increase trust. Second, the Information Commission (ICO) regulates the use of personal data in political campaigning, which is the basis of microtargeting. Third, the Advertising Standards Authority’s CAP Code places requirements on advertisers in relation to misleading advertising, but it excludes political advertising from its voluntary remit. This is a questionable gap, but at present it is up to it, not any regulator, to close it. Fourth, the response to state-backed interference campaigns is principally a matter for Government, including through police and security services. Finally, the service providers themselves inevitably remain the first line: their terms of service, advertising policies, and integrity teams govern far more conduct, far faster, than any regulator can.

Conclusions

This advice note has mapped the landscape of information-based threats to elections, offering a high-level conspectus of the literature review we commissioned on this topic. In offering this overview, we have sought to recognise the real harms and risks posed by varieties of electoral disinformation, while also grappling with how little is actually known about their magnitude. We then sought to spell out Ofcom’s role within the broader landscape, stressing what measures it can and should take to combat this threat while upholding freedom of expression. But we also emphasised that Ofcom’s powers are highly limited, and that it is only one actor in a broader division of labour between regulators, Government, and industry.

Our advice to Ofcom (which includes but goes beyond the points raised in the previous section) is as follows:

- Ofcom should reiterate, well in advance of the next UK General Election, its public statement of its role, and the limits of its role, in relation to electoral mis- and disinformation. It is imperative to educate the public on what the Online Safety Act empowers it to do, and what it doesn’t.
- We advise Ofcom to ensure that regulated services are complying with their duties to take action against illegal content, especially priority illegal content. Regulated services have duties to protect their users from illegal harms involving foreign interference, harassment, threats, non-consensual intimate image abuse, and false and threatening communications, among other crimes. We recognise that the foreign interference offence involves a high bar meaning that it may be particularly difficult for providers to identify online content as amounting to the offence due to challenges of attributing activity to a foreign power, and much of the evidence on threat assessments is not public.
- We advise Ofcom to maintain its schedule to finalise its guidance regarding Category 1 services’ additional duties on terms of service. Category 1 services have a duty to ensure that their terms of service are enforced consistently. Many services have voluntarily adopted terms of service against electoral misinformation, fake or impersonating accounts,

¹⁵ [Elections Act 2022, Part 6](#) (digital imprints). Note that the Commission has no role in adjudicating the truth of campaign claims.

coordinated inauthentic behaviour, and deceptively manipulated media; holding services accountable for their terms of service is thus an important mechanism in limiting the spread of electoral disinformation.

- We advise Ofcom to maintain its schedule to finalise its guidance regarding Category 1 services' additional duties on content of democratic importance and journalistic content, and on conducting freedom of expression impact assessments. This will help mitigate against overenforcement and reduce the number of false positives.
- Ofcom should continue its media literacy work in educating UK users on how to spot unreliable sources and falsehoods in the context of elections. Because Ofcom lacks powers to compel platforms to implement its media literacy recommendations, it is imperative for platforms to implement them voluntarily. Ofcom should leverage its supervisory relationships to urge platforms to implement recommendations from its voluntary guidance on best practices. And Ofcom should use its transparency powers to compel platforms to disclose information on their own efforts to promote media literacy.
- Ofcom should support the Electoral Commission in discharging its duty to promote public awareness of elections. Ofcom should seek to work with the Electoral Commission to promote the visibility and findability of authoritative electoral information across the languages spoken in the UK, support early communication about electoral processes in those languages, and work with schools to prepare first-time voters. This is particularly significant given the extension of the franchise to sixteen- and seventeen-year-olds.
- Ofcom should encourage and commission research into the gaps this review identifies. These should focus on (1) the most effective media literacy strategies for people under 18; (2) understanding how people are using generative AI to seek out political information, the ways in which mis/disinformation can emerge in generative AI, and how generative AI use is influencing voting beliefs; and (3) what might foster people's trust in authoritative information during an election. Ofcom should also continue to support improved researcher access to platform data, without which the evidence base cannot mature.
- Ofcom should plan to collect data on, and evaluate the effectiveness of, measures taken under the illegal content duties in respect of the foreign interference and other relevant offences. Ofcom's transparency powers should be directed at the areas where evidence of malicious conduct is clearest: coordinated inauthentic behaviour, fabricated and synthetic accounts, provenance-obscuring networks, and spreading false content about the logistics of elections.
- Ofcom should take seriously the risk that disinformation targeting politicians as a form of harassment is deterring people from running for office, especially women and those from ethnic minority backgrounds. Ofcom's own research in this area shows that this is a serious problem.¹⁶ Ofcom should robustly use its enforcement powers in relation to suspected non-compliance with duties relating to existing priority offences. It should also use the recent [call for evidence](#) on a safer life online for women and girls to strengthen the existing evidence base.

¹⁶ Ofcom, 2025. [Experiences of online hate and abuse among women in politics](#).