External Undermining

Methodological memorandum — Purpose, scope and definition, dossier architecture, the severity/certainty distinction, and limitations

1. Introduction and problem statement

The External Undermining dashboard is a citizen-facing transparency instrument that documents, per EU member state and at institutional level, attempts by non-EU actors to destabilise democratic processes within the Union, influence public opinion, and erode trust in politics. It is the third instrument in the series that also comprises the EU Governance Assessment (EUGA) and the Lobby Dashboard.

Where the EUGA measures governance outcomes at member state level, and the Lobby Dashboard addresses influence exerted from within the Union's own legislative process, this instrument addresses pressure applied from outside it. The three are complementary: internal institutional erosion, internal but under-transparent influence, and external interference are distinct phenomena that require distinct evidentiary treatment. Conflating them would obscure rather than clarify.

The problem this memorandum addresses is that external interference is comparatively well documented but poorly accessible. The evidence exists — in national intelligence assessments, in the annual FIMI reporting of the European External Action Service, in the work of specialised research institutes, and in investigative journalism — but it is dispersed across languages, formats and institutional silos. No public instrument brings it together per member state in a form a citizen can search, compare and verify.

2. Need and purpose

2.1 The problem: documented, but not accessible

A citizen who wishes to know whether, and how, foreign actors have sought to influence democratic processes in their own country faces a practical obstacle rather than an absence of evidence. National security services publish annual reviews, but in the national language and rarely in a comparative frame. The EEAS publishes FIMI threat reports at Union level, but largely in aggregate terms. Investigative journalism documents individual cases, but episodically and without a common vocabulary.

The consequence is that the picture available to the public is simultaneously alarming in tone and thin in specifics. That combination is itself corrosive: it invites both complacency and over-attribution, and it leaves citizens without the means to distinguish a well-evidenced case from a contested allegation.

2.2 What existing instruments do not offer

Whether a comparable instrument already existed was investigated explicitly during development. Several established instruments cover part of the ground, but none provides a per-member-state dossier structure with an explicit, visible evidentiary grading.

Existing instrument	What it does provide	What it does not provide (and this dashboard does)
EEAS FIMI threat reports	Authoritative annual EU-level analysis of foreign information manipulation, with incident counts and actor characterisation.	Per-member-state dossiers; case-level detail for most countries is reported only in aggregate.

Existing instrument	What it does provide	What it does not provide (and this dashboard does)
National security service reports	Detailed, high-authority findings for the reporting country, often naming actors and methods.	Comparability across member states; published in national languages, on differing cycles and to differing standards.
EUvsDisinfo database	A large, searchable corpus of individual pro-Kremlin disinformation cases.	Structured attribution of severity and certainty per case; organisation by member state as a dossier rather than by narrative.
Specialised research institutes (DFRLab, ISD, CSD and others)	Rigorous case investigations with technical attribution evidence.	Systematic coverage of all member states, including those where nothing was found.

Table 1. Existing instruments and the gap the External Undermining dashboard addresses.

2.3 Empirical grounding: what the twenty-nine dossiers show

The case for such an instrument is not only theoretical; it was confirmed empirically during compilation. Coding all 27 member states plus the EU institutional level and NATO produced 100 documented incidents and three findings that were not visible from any single source.

First, interference is selective rather than universal: effort concentrates where the expected return is highest, and post-election analysis in Denmark concluded explicitly that the country had been treated as a low-return target. Second, the dominance of Russia in the aggregate figures (roughly four in five attributed incidents) partly measures which actor is most actively tracked and reported, not only which actor is most active. Third, the actor picture is broadening in ways aggregate reporting obscures: China disproportionately in Ireland, Belarus as a separately attributed actor in Lithuania, and a private intelligence firm in Slovenia.

Each of these observations depends on being able to compare member states on a common evidentiary basis — which is precisely what the dispersed source landscape does not permit.

3. Scope and definition

3.1 Working definition

For this dashboard, external undermining is defined as: activities by non-EU actors that, financially or otherwise, support parties or campaigns with a view to destabilising the democratic process within the EU, influencing public opinion, and undermining trust in politics.

Two elements of this definition do substantive work. The actor must be external to the Union, and the activity must bear on democratic process or public trust rather than on state interests generally. Both boundaries exclude conduct that is serious in its own right but belongs to a different instrument.

3.2 What falls within scope

- Disinformation campaigns directed at elections, referendums or public opinion generally.
- Covert or non-transparent financing of parties, campaigns or individual politicians by foreign actors.
- Network and infrastructure operations — imitation media sites, bot networks, proxy outlets — with a demonstrable political purpose.
- Cyber operations with an electoral or society-disrupting component.
- Direct narrative interventions by state actors themselves, including public statements by a foreign intelligence service about a domestic party.

3.3 What is deliberately excluded, and why

Four categories of conduct are excluded by design. In each case the exclusion is a boundary of this instrument, not a judgement that the conduct is unimportant.

Excluded category	Illustrative cases encountered	Reason for exclusion
Purely domestic disinformation	The fabricated Tisza election platform in Hungary; pro-Kremlin statements by Italian politicians; the Austrian 'Ibiza affair'.	No external actor. The Ibiza sting was a journalistic operation, not a foreign actor's action.
Intra-EU interference	An allegation that Poland interfered in the Hungarian election campaign.	The definition specifies non-EU actors. Excluding it is a scope boundary, not an assessment of gravity.
International organisations without a collective defence function	The OPCW, headquartered in The Hague.	Being hosted on EU territory does not create the functional link to member-state security that justifies NATO's inclusion.
Physical hybrid threats without a political link	GPS jamming and comparable incidents.	Coded to the affected member state, and only where a source itself draws the link to political influence.

Table 2. Categories excluded from the incident count, with the reasoning applied.

4. Methodological approach

4.1 Design principles

- One member state, one dossier, four questions: how large, who was targeted, which actor, and to what end.
- Separation of finding and interpretation: what a source states is reported as such; the dashboard adds no attribution the source does not make.
- A consistent base structure across all dossiers, with deliberate and explained deviations where the evidentiary situation genuinely differs.
- No invented precision: where attribution is contested or absent, this is displayed rather than resolved.
- Symmetry: the same research effort and the same evidentiary threshold apply to every member state, irrespective of size or region.

4.2 Building blocks per dossier page

Each dossier is built from a fixed set of components, so that member states remain comparable even where the underlying evidence differs sharply in depth.

Building block	What it shows	Basis and constraint
Scale and timeline	Incident count per year and a chronological account of each case.	Derived from the incident database; every entry carries a source and date.
Targets	The election, institution, party or population group affected.	Recorded as the source frames it, not inferred from context.
Actors	The state or network to which activity is attributed, with certainty per case.	Joint and unattributed cases are retained as such rather than assigned to a single state.

Building block	What it shows	Basis and constraint
Aims	What the source describes as the intended effect.	Rephrased from the source; no independent interpretation of motive is added.
Sources and methodology	Full source list with certainty grading, plus dossier-specific caveats.	Each source is linked; caveats are written per dossier, not standardised.

Table 3. Building blocks of a dossier page.

4.3 Why dossiers rather than a composite index

The EUGA is built as a composite score because it rests on annual expert indices with full country coverage. That precondition does not hold here. External interference is covert, varies widely in form, and the available sources cover member states unevenly — a function of national reporting practice as much as of underlying activity.

A composite score built on such a base would suggest a precision the data cannot support, and would convert a reporting artefact into an apparent measurement. The dossier approach, following the Lobby Dashboard, keeps each incident traceable to its source and makes the unevenness of the evidence visible rather than averaging it away.

4.4 Severity and certainty as two independent scales

A recurring analytical finding during compilation is that a case being striking and a case being well-substantiated are different properties, and that conflating them distorts the picture. Severity and certainty are therefore coded as two independent variables which are never combined into a single rating.

The Slovak dossier illustrates why this matters. The two viral deepfakes released before the 2023 election carry a high severity rating but a low certainty rating, because their creators were never identified. The most certain case in that dossier — a public statement by Russia's own foreign intelligence service — carries only a medium severity rating. A single blended score would have inverted the relationship between prominence and evidentiary strength.

4.5 Source hierarchy and certainty rubric

The certainty rating reflects the strength of the attribution, not the seriousness of the conduct. It is applied per incident, using a fixed three-point rubric consistent with the source certainty conventions used elsewhere in the project.

Certainty	Basis	Typical source
High	Attribution by a state body with investigative authority, an EU institution, or confirmation across multiple independent sources.	National security service; EEAS; judicial verdict.
Medium	Attribution by a specialised institution with relevant technical or analytical competence.	Research institute; cybersecurity firm; academic publication.
Low	Attribution resting on a single journalistic source, or a claim not independently substantiated.	Single press report; unconfirmed official allegation.

Table 4. Certainty rubric applied to each coded incident.

4.6 Treatment of data gaps and the 'searched, none found' status

Where reliable data are unavailable, this is displayed rather than filled. The dashboard distinguishes explicitly between a member state that has not been examined and one that has been examined without a qualifying incident being found — a distinction that most comparable instruments collapse.

Luxembourg is the sole instance of the second category. Three search passes in two languages, including searches targeted at the named networks documented elsewhere in the dataset, produced only generic threat-awareness statements. Rather than omitting the country or admitting a weaker source to fill the gap, the dossier records the search process itself and is marked distinctly on the map. An absence of evidence is thereby prevented from reading as an absence of enquiry.

4.7 Deliberate deviations from the standard structure

Two entries depart from the member-state template because the unit of analysis genuinely differs.

The EU institutional level exists for cases in which a source frames the target as an EU institution or a Union-wide process — the European Parliament elections taken as a whole, for instance — rather than as a national matter. Without it, such cases would either be attributed arbitrarily to one member state or lost between dossiers.

NATO is included as a separate entry, while other international organisations situated on EU territory are not. The distinction rests on function rather than location: NATO collectively defends member-state territory, so its erosion bears directly on the security of the states this dashboard covers. The OPCW, by contrast, is hosted in the Union without that structural relationship. The NATO dossier is deliberately the thinnest in the set, because the available sources describe the threat only in aggregate; that limitation is stated within the dossier itself.

5. The dashboard: presentation layer

5.1 Why a presentation layer is necessary

The incident database is auditable but not readable. A spreadsheet of 100 coded incidents across 29 entries answers questions only for a user who already knows what to ask. The dashboard exists to make the same material navigable for a reader who does not — by country, by actor, and by strength of evidence.

5.2 Landing page: map, actor summary and tiles

The landing page combines a hexagonal cartogram, an aggregated actor summary, and an alphabetical tile list. The actor summary is calculated automatically from the certainty-tagged data held in each dossier rather than maintained separately, which prevents the headline figures from drifting out of step with the underlying records. Joint attributions are reported in their own category rather than distributed across the states named, and unattributed incidents are shown as such.

5.3 Country dossiers

Each dossier opens on a fixed five-tab structure corresponding to the building blocks in Table 3. Where a tab has no content — as with the targets and aims tabs in the Luxembourg dossier — an explanatory statement is displayed rather than an empty panel, so that the absence is legible as a finding rather than as a fault.

5.4 Cartogram conventions

The map is a schematic grid cartogram with equal cell size per entry, not a representation of borders or territory. Equal cells prevent small member states from disappearing beside large ones while preserving relative position. Switzerland is shown as a non-clickable outline so that the gap in the grid reads as a deliberate geographic feature rather than a missing member state.

Colour intensity encodes the number of documented incidents, not their severity. A state with several minor incidents would otherwise appear graver than a state with one serious one. Severity and certainty are available per incident inside the dossier, where they can be read with their qualifications attached.

5.5 The teaser

A half-A4 landscape teaser accompanies the article, carrying the map, the actor summary and a link through to the full dashboard. It follows the same format and house style as the EUGA teaser so that the three instruments present consistently within the publication.

6. Publication and update cycle

The recommended cycle is annual, following publication of the EEAS annual report on FIMI threats, which is the single most comprehensive source in the set and typically appears in March. National security service reviews, which appear on differing national cycles, are incorporated as they are published.

New incidents are added to the incident database first; the dashboard and teaser figures follow from it. Because the actor summary is computed rather than maintained by hand, routine additions do not require a design revision.

7. Source overview

The table below sets out the principal categories of source used across the twenty-nine dossiers. Every individual incident carries its own named source, date and link in the incident database.

Category	Examples of sources used
National intelligence and security services	BfV (Germany), VIGINUM (France), AIVD/MIVD/NCTV (Netherlands), SUPO (Finland), SÄPO (Sweden), PET (Denmark), KAPO (Estonia), VSD (Lithuania), VDD (Latvia), BIS (Czechia), SOVA (Slovenia), SIS (Portugal), VSSE (Belgium), DSN (Austria).
EU institutions	European External Action Service FIMI threat reports; EUvsDisinfo; European Parliament resolutions and parliamentary questions; European Commission sanctions decisions.
Specialised research institutes	Atlantic Council DFRLab; Institute for Strategic Dialogue; Centre for the Study of Democracy (Sofia); NATO StratCom Centre of Excellence; Real Instituto Elcano; SCEEUS.
Investigative journalism and civil society	VSquare (#ESPIOMATS); Denik N; DR; Balkan Insight; BG Elves; MediaLab (Iscte); EDMO and its national hubs.
Judicial and parliamentary findings	Vienna regional court verdict in the Ott case; Romanian CSAT declassified documents; Italian COPASIR annual report.

Table 5. Principal source categories used across the twenty-nine dossiers.

8. Why this dashboard serves an essential monitoring function

Article 2 TEU names democracy, the rule of law and respect for human rights as the values on which the Union rests. Democracy in that sense presupposes that citizens can form political judgements freely and on an informed basis, without covert influence by actors outside their own democratic community. External undermining attacks that precondition directly. This chapter states why the combination described in chapters 1 to 7 fulfils a role that no existing instrument performs.

- **It closes the gap between available evidence and citizen accessibility.** The material exists across intelligence reporting, EU institutions and research organisations, but not in a form the public can search by country and compare on a common basis.

- **It is organised by member state rather than by narrative or campaign.** Existing instruments catalogue disinformation cases or describe networks; neither answers the question a citizen is most likely to ask, which is what has been documented in their own country.
- **It makes evidentiary strength part of the output.** By separating severity from certainty and displaying both, the dashboard allows a reader to distinguish a court-confirmed case from a contested allegation — a distinction that aggregate threat reporting tends to flatten.
- **It records negative findings.** Documented absence, as in Luxembourg and to a large extent Denmark, is treated as a result. This is what lends the positive cases their credibility, and it guards against the assumption that every member state is targeted equally.
- **It makes its own choices auditable.** Every scope boundary, coding rule and presentational convention is set out in this memorandum, and every incident is traceable to a named source in the accompanying database.

9. Consolidated limitations of the dashboard

The following limitations apply to the instrument as a whole. Dossier-specific caveats are stated within the dossiers themselves and are not repeated here.

- **Publication and detection bias.** Member states with an active national reporting practice show systematically more incidents than those without. The count reflects the capacity and willingness of sources to publish as much as the underlying level of activity.
- **Concentration on a single actor.** Russia is the attributed actor in roughly four of five coded incidents. This may reflect the real distribution of activity, or the distribution of Western investigative attention. The dashboard does not adjudicate between the two.
- **Point-in-time snapshot.** Each dossier reflects the state of public documentation at its coding date. Ongoing investigations, particularly the recent Slovenian and Hungarian cases, may materially change.
- **Contested attributions remain contested.** Where sources conflict, both readings are presented as the sources present them. The dashboard does not resolve disputes on its own authority.
- **Thin dossiers are not evidence of low risk.** A low incident count indicates limited public documentation, not necessarily limited exposure.
- **No causal claim.** The dashboard documents attempted interference and its documented targets. It does not establish that any campaign changed an electoral outcome, and no such inference should be drawn from incident counts.
- **Manual compilation.** The instrument is a manually coded set of dossiers, not an automated feed. Coverage depends on continued curation.
- **Schematic cartography.** The grid map makes no claim about borders, territory or sovereignty; cell positions serve legibility only.

10. Status of this document

This memorandum is the methodological account of the External Undermining dashboard and is intended to be read independently of it. It sets out the definition applied, the boundaries drawn, the coding rules used and the limitations accepted, so that a reader can assess the instrument's claims against the choices that produced them.

The dashboard is not peer-reviewed academic research and does not present itself as such. It is a documentation instrument built on public sources, with its evidentiary limits stated rather than concealed. Where the evidence

is thin, the dashboard says so; where attribution is disputed, it leaves the dispute visible. That is the standard against which it should be judged.

Compiled by Aantekeningen in de marge, with extensive editorial and visual input from Claude, Anthropic's AI assistant.