

## Appendix J: An illustrative evidence infrastructure

The main text mentions, as one option in the French context, a standing observatory of the information environment. This appendix sketches such a body and weighs the case for and against it; the diagnosis it builds on is France’s, and we cannot yet be sure it holds in other countries. Since our fieldwork, the United Kingdom has announced a National Centre for Information Defence, whose mandate, powers, and resources remain unspecified (Burnham, 2026). As announced, it would detect, attribute, and disrupt hostile-state campaigns, close to VIGINUM’s remit; the measurement and translation functions sketched below could complement it. It is not a policy program: reviews of what works already exist in English (Altay, 2026; Kozyreva et al., 2024), alongside evidence-based guides and briefs that distill, for decision-makers, what science says about the policies to pursue (e.g., Bateman & Jackson, 2024; Lewandowsky et al., 2020), including recently in French (e.g., Darcy et al., 2025; DITP & VIGINUM, 2026). Those resources say what to do; the question here is how the response system itself could work better.

Two of this study’s findings give the proposal its shape, speculative though it remains. Practitioners lack shared measures: 55% rate the absence of common indicators a serious obstacle, and institutions count what is easy to count (campaigns detected, content removed), so the habit of watching only the supply side keeps confirming itself. A second gap runs alongside the first: practitioners trust research without reading it, with 71% trusting peer-reviewed work, 36% keeping up with it, and 88% reading about disinformation in French while the evidence appears mainly in English. An observatory could supply what is missing on both counts: a unit to measure what current tools leave invisible, and a unit to translate existing research for those who act on it.

Its first task could be to get the institutions named here to agree on a handful of common indicators, and the precedents warn how hard that is. The European Monitoring Centre for Drugs and Drug Addiction (EMCDDA, now the European Union Drugs Agency, EUDA), the closest case of credible measurement in a politicized field, needed years to harmonize five indicators it could anchor partly in biological data; “exposure to disinformation” has no agreed definition.

Four kinds of work could fall to the measurement unit. It could begin with a recurring panel survey tracking what current tools miss: exposure, beliefs, trust, and how they shift over years. The Institut national de la statistique et des études économiques (INSEE) could field and hold it under statistical secrecy, releasing only anonymous aggregates, so that no state-linked body ever holds identified citizens’ beliefs; the long series matters, since nearly all respondents expect the threat to worsen and only 13% expect the response to keep pace (Finding 5). It could add the supply-side measures current instruments miss: how many journalists cover local affairs in each territory, how news is financed, how trust in it evolves. It could also draw on the platform data that vetted researchers can now request under European Union (EU) regulation (Article 40 of the Digital Services Act, whose data-access mechanism became operational in 2025), using it to see how disinformation actually circulates: which accounts and networks amplify it, how fast, and to whom. That access could help, but it is granted project by project for a limited period, and platforms have restricted or withdrawn researcher access before, so a ten-year series cannot rest on it. Finally, it could test whether the interventions the field already funds actually work, running randomized or quasi-experimental trials of them (of a media-literacy campaign, say, or a fact-checking program) and standardized reviews after each major incident. One feature could cut across all four: demand-side measures (institutional trust, isolation) would appear in the same instrument, and on the

same schedule, as the foreign-campaign counts. Because diagnoses tend to follow whatever the instruments measure, the instrument itself would have to correct the selective diagnosis of Finding 2.

The goal of the knowledge translation unit would be to bring existing research to the people who can act on it. The research they need largely exists; the trouble is that it reaches them in English, behind paywalls, in formats remote from the decisions they have to make. The evidence-use literature shows both how to close that gap and why its effect sizes stay modest. Sending people research does little on its own: a large randomized trial, which sent evidence-based materials to hundreds of schools, changed neither teachers' use of research nor pupil outcomes (Lord et al., 2017); access helps mainly when it is paired with motivation and opportunity to use it (Langer et al., 2016). Both are present here: practitioners already trust research (71%) and seek operational answers, and the obstacles they report are practical ones of time, access, and incentive. A portal alone would go unused.

Instead, the unit could deliver four kinds of product. It could write, first, plain-language French syntheses, open access, summarizing what research says on the questions practitioners face repeatedly (does fact-checking work, how much exposure is there, do literacy programs transfer); none yet exists systematically in French, the language 88% of respondents read. It could tailor, second, role-specific products (a legal memo for the regulator, a threat note for the officer, a classroom resource for the educator), on a known schedule, with a channel carrying practitioners' questions back to researchers. It could assign, third, knowledge brokers selectively: a large randomized trial of knowledge brokering found brokers helpful only where research culture was weak (Dobbins et al., 2009), which in our sample points to the communication/counter-influence and public-protection roles that find research least applicable (Appendix D.3). And it could give, fourth, regular indicator feedback, which yields small gains (a median of 2.7 percentage points in the latest Cochrane review) that grow larger when the feedback compares performance to peers and comes with a concrete action plan (Ivers et al., 2026).

The same unit could also study translation itself. Research-policy initiatives almost never check whether they work (6% of some 2,300 in one census; Oliver et al., 2022); a unit that tracked its own uptake from the start could correct course early and produce some of the first systematic evidence on how translation works in this field.

The precedents show what could threaten such a body: in these institutions, survival is dictated mainly by funding. The endowed Education Endowment Foundation has run trials for over a decade, while the What Works Centre for Wellbeing closed in 2024 once its grants lapsed and no ministry owned its subject. The Stanford Internet Observatory's election work, which tracked false and misleading claims about the 2020 and 2022 elections, drew lawsuits, subpoenas, and congressional investigations that cost millions of dollars in legal fees; by mid-2024 its leadership had left and the election project had ended (Bond, 2024). Europe's media observatories, for their part, depend on renewable EU grants, which makes a long time-series hard to sustain. We can draw on these precedents to set three rules. Fund the body once (through an endowment or a multi-year statutory commitment) rather than renegotiate annually; give it public-law status, so that the state itself bears the legal defense; and say plainly what it will never do: no content flagging, no truth adjudication, no election-time operations, no data sharing with security services. That last rule speaks to a finding of this study: practitioners are inhibited less by the ethical dilemmas they actually face, which 21% encounter often, than by the censorship accusations they anticipate, which 63% rate a serious obstacle. A statute that made the accusation false could lower the cost of defending legitimate action and make unethical action both less likely and harder to carry out.

One of the main objections would be that the body merely duplicates what France already has: a detection service, a regulator that runs surveys, a European hub, and a statistical institute. Yet none is placed to pull these strands together. VIGINUM (Vigilance et protection contre les ingérences numériques étrangères) is intelligence-adjacent, so a panel of citizens' beliefs in its hands would not be trusted; the Autorité de régulation de la communication audiovisuelle et numérique (ARCOM) cannot credibly evaluate the policies it enforces; the European hub has no stable funding; INSEE commands the public's

trust and the survey craft but holds no mandate over the information environment and performs no translation. More fundamentally, this very dispersion is the problem the body could address: each ministry treats “its own” disinformation (the Health Ministry vaccine falsehoods, the Interior Ministry electoral interference, Foreign Affairs and VIGINUM foreign campaigns, ARCOM platform compliance) with its own instruments and definitions, so no one holds the shared picture and no institution leads (Darcy, 2026b). Respondents confirm it: 54% rate unclear responsibilities a serious obstacle.

One of the main vulnerabilities would be political, and it follows directly from that dispersion. The dominant story places the threat abroad. That story is convenient, because acting on causes at home, such as falling trust in institutions and a weakened local press, would implicate domestic choices (Darcy, 2026b); a government that gains from blaming outsiders has little incentive to fund the instrument that could correct it. Because that incentive is missing, the body could not arrive by decree; it could only be built in stages: first federate what already exists (a stable funding line for the existing hub, a recurring statistical-institute module, a small council commissioning evaluations), with the common indicators as the test, and institutionalize only if that work converges.

Two design risks would remain even then. Indicators that become targets get gamed (Bevan & Hood, 2006; Campbell, 1979), so the suite would have to rotate and carry its uncertainty openly; and because published measures of belief can be turned against the groups they describe, results should never be ranked by population group. Other questions stay open too: the body’s resilience to political alternation, its independence, and the protection of the data it would hold. None of this makes the observatory a sure solution. It is one way, among others, of acting on what this study documents: a field asked to fight a phenomenon it cannot measure, with evidence it cannot read.