

Appendix A: Participants, sampling, and methods

Target population

The target population comprised practitioners who, by virtue of an official position in a public-sector, private-sector, or non-profit organization, hold responsibility for preventing, monitoring, analyzing, regulating, or responding to false, misleading, manipulated, or hostile informational content circulating in public communication channels. Eligibility required either (i) an explicit institutional mandate in communication or counter-influence; normative frameworks, regulation, or public policy; detection, monitoring, and analysis; public protection, journalism, or education; or cybersecurity and technical protection, including coordination and partnership functions within institutions formally tasked with counter-disinformation, or (ii) demonstrable functional responsibility for disinformation response within a broader portfolio (e.g., officers overseeing operations targeted by hostile influence campaigns, or senior communications staff supporting high-level political principals confronted with recurrent misinformation episodes).

Academic researchers were excluded by design. Their beliefs and policy preferences have been documented in prior surveys, notably by Altay and colleagues; the present study instead focuses on field actors who make and implement decisions under administrative, political, legal, and, in some settings, security constraints.

Sampling strategy

A probabilistic sampling design was not feasible because no comprehensive institutional mapping of “disinformation practitioners” exists in France. Mandates are distributed across organizations, job titles are heterogeneous, and responsibilities are often embedded within broader roles, preventing construction of a complete sampling frame. We therefore implemented an identification strategy aimed at maximizing coverage of the reachable practitioner ecosystem. Eligible participants were identified using public organizational charts and official documents, participation in conferences and working groups, interagency initiatives, professional networks and prior contacts, and snowball referrals. Because inclusion depended on identification, institutional authorization, and willingness to participate, the sample should not be interpreted as a random draw from the universe of all French practitioners working on disinformation.

Recruitment and data collection

Recruitment and data collection proceeded in two waves, from September 2025 to February 2026. In Wave 1, identified practitioners were contacted directly by email and, when feasible, in person, including ministerial cabinets and agencies formally mandated to counter foreign information manipulation such as intelligence services or VIGINUM (the French government service for vigilance and protection against foreign digital interference). Given heterogeneous confidentiality and authorization constraints (especially in defense and intelligence settings), a single standardized administration mode could not be imposed. Depending on institutional rules, respondents completed the questionnaire either with a

researcher present in person, or autonomously via a secure link, with the research team reachable at any time for clarification. After applying inclusion criteria (valid consent and confirmed professional role), Wave 1 contributed 80% of the analytical sample.

Wave 2 was implemented to improve coverage of practitioner segments that remained underrepresented after direct outreach, particularly private for-profit and private non-profit organizations. An anonymous Qualtrics link was therefore disseminated through two targeted channels: a professional community dedicated to media and digital literacy education and a newsletter circulated within a professional network of people engaged in counter-disinformation, coordinated by a member of the French armed forces. After applying the same inclusion criteria, this pathway contributed the remaining 20% of the final sample.

Sample composition

The final analytical sample ($N = 106$) spans six organizational types: military and intelligence services (33%), regulatory and normative bodies (29%), senior government and political offices (21%), private for-profit organizations (7%), private non-profit organizations (6%), and journalism and media (4%). Detailed institutional composition is described in the main text.




Figure A1. Sample composition and exposure to disinformation (% , N varies by panel: 102–106). Professional situation, educational background, main professional role, and whether respondents' institutions had been directly targeted by disinformation. Waves 1 and 2 combined.

The survey recorded no personal identifiers, so individual respondents cannot be identified from their answers. To preserve confidentiality, particularly for diplomats, members of presidential or ministerial cabinets, secret-service agencies, or the military, we collected no demographic data beyond professional role, type of organization, and educational background (no gender, age, or ethnicity). This also motivated

certain category constructions (e.g., grouping “military and military agencies and secret services” or “diplomacy with high administration and politicians”) to prevent any possibility of re-identification from cross-tabulated data, which is a common practice for studies dealing with members of intelligence agencies.

Inclusion criteria and analytic sample

The survey received 123 consented responses. We excluded 17 cases that did not complete the initial demographic block (i.e., did not answer the professional-role screening question, Q7). Of these, 71% ($n = 12$) dropped out immediately after consent (survey progress = 5%), and one completed most of the survey (78 of 92 items) and provided demographic information but did not answer the screening item (Q7). Four additional cases reached intermediate progress (33–95%) but did not provide a valid professional-role response. Exclusions were disproportionately drawn from later recruitment batches, which may reflect cohort-specific response-rate variation. Among the five excluded cases with partial demographic data, three reported a sector (non-profit, $n = 1$; other, $n = 2$) and two reported only their educational background, suggesting that attrition was not concentrated in a single sector. The remaining 12 provided no demographic information. We treated completion of the professional-role item as a minimal completeness filter for inclusion in the analytic sample. By contrast, no respondent selected the screening option stating that their main activity was unrelated to disinformation, so no participant had to be excluded on that ground. This is informative in itself because it suggests that the targeted recruitment strategy reached the intended practitioner population rather than attracting a visible tail of topic-interested outsiders. The analytic sample therefore consists of $N = 106$ respondents.

Missing-data treatment

Given the small and specialized population, analyses rely on pairwise deletion, or available-case analysis, which retains all valid responses for each item rather than discarding an entire respondent because of a single missing value. In expert samples, listwise deletion may inadvertently privilege participants with greater time or institutional latitude to complete every question, thereby introducing sectoral distortions. Item-level denominators consequently vary across blocks: perceived impact (Q9, $N = 101$ – 102), reasons for belief (Q12, $N = 99$), mechanisms of spread (Q13, $N = 96$), obstacles (Q16, $N = 93$), research engagement (Q22, $N = 91$), and research obstacles (Q23, $N = 87$ – 88). Within most blocks, all items share the same denominator, meaning that missingness is block-level rather than item-level. The partial exceptions are Q8 (relative importance, $N = 101$ – 102), Q9 (perceived impact, $N = 101$ – 102), and Q23 (research obstacles), where minor item-level variation occurs; for Q23, the first five items each have $N = 87$ while the sixth has $N = 88$, a discrepancy of a single respondent that does not materially affect estimates. For all other multi-item blocks (Q10, Q12, Q13, Q15, Q16, Q19, Q22), pairwise and listwise deletion yield identical point estimates because respondents who answered any item within a block answered all items in that block.

To assess whether missingness was related to respondent characteristics, we conducted 48 χ^2 tests of independence crossing the 16 closed-ended question blocks with three grouping variables (organizational sector Q3, educational background Q5, and professional role Q7). After Benjamini–Hochberg correction (Benjamini & Hochberg, 1995), only one test reached significance (Q8 missingness $\times$ educational background, $p_{BH} < .001$), driven entirely by a single respondent in the mathematics/statistics category ($n = 1$, 100% missingness)—a cell-size artifact rather than a substantive pattern. Five tests were significant before correction: Q8 $\times$ Q5 ($\chi^2 = 35.27$, $df = 7$, $p < .001$), Q2 $\times$ Q5 ($\chi^2 = 20.88$, $df = 7$, $p = .004$), Q22 $\times$ Q5 ($\chi^2 = 16.67$, $df = 7$, $p = .020$), Q8 $\times$ Q7 ($\chi^2 = 13.06$, $df = 5$, $p = .023$), and Q16 $\times$ Q5 ($\chi^2 = 15.31$, $df = 7$, $p = .032$). All five involved minimum expected cell counts below 5, which inflates the χ^2 statistic and makes the

asymptotic approximation unreliable; Fisher exact tests returned non-significant p-values for three of the five. These uncorrected associations are concentrated in Q5 (educational background) and trace back to very small subgroups (mathematics/statistics $n = 1$, natural sciences $n = 4$, cognitive science $n = 2$), where a single missing respondent can produce a large percentage shift. Although no robust evidence of a missing-completely-at-random (MCAR) violation was found, the educational-background dimension warrants interpretive caution for blocks with higher missingness rates (Q16, Q22, Q23).

At the block level, representative tests confirm the absence of large-scale sectoral confounding: for example, a χ^2 test of Q23 (the final question) missingness by sector is non-significant ($\chi^2 = 8.30$, $df = 6$, $p = .217$). Missingness in Q23 is nonetheless uneven in absolute terms across organizational sectors: it is highest in private non-profit organizations (50.0%, 3 of 6) and private for-profit organizations (28.6%, 2 of 7), and lower in military/intelligence services (8.6%, 3 of 35) and among journalists (0%, 0 of 4). Under listwise deletion, these differences would mechanically remove disproportionate shares of certain sectors and bias aggregate estimates. Pairwise deletion preserves information from earlier sections while making variation in denominators explicit.