

ANSWERED BEFORE THE COMMITTEE ASKS — INCLUDING THE QUESTIONS WE CANNOT CL

AI-DSM

ETHICS DOSSIER

AI-DSM Ethics Dossier for the Behavioural Study of Grown Systems



This study deliberately creates flawed model checkpoints, measures relational and agentic failures that have already reached courts, and publishes named findings about commercial products. Every one of those is defensible only with an architecture stated in advance. This dossier states it, names what is still missing, and reads itself as a hostile committee would.

Contained induction

Independent veto

No user ever in the loop

Machine-welfare caution

Register of open i

Stephane van der Aa

15 September 2026

stepvda.substack.com

Contents

Before the committee asks.....	4
The four statuses, and what each means.....	4
Which framework governs which part of this programme.....	4
Six positions this dossier takes, and defends.....	5
The study in one page, for a reader coming to it cold.....	5
Part 1 — Why this dossier exists and how to read it.....	6
1.1 What a committee actually does with a submission.....	6
1.2 How this document is organised.....	6
1.3 What this dossier is, and what it is not.....	6
Part 2 — The legal and regulatory frame.....	7
2.1 What kind of research this is, legally.....	8
2.2 Which committee reviews a study whose subjects are checkpoints.....	8
2.3 The sponsor, and whether that legal person exists.....	9
2.4 The EU AI Act, the research exemption, and the 2026 Digital Omnibus.....	9
The research exemption, and what it does not cover.....	10
The category trap, and why it is a containment rule.....	10
GPAI obligations and the Code of Practice.....	10
What the 2026 Digital Omnibus changed.....	10
2.5 Data protection law, in outline.....	10
2.6 Dual use, export control and provider terms.....	11
Does creating deliberately misaligned checkpoints engage a dual-use or export-control rule?.....	11
Provider terms of service, and what the pilot did.....	11
2.7 Liability, insurance and indemnity.....	11
Part 3 — Scientific validity as an ethical requirement.....	12
3.1 The evidence base moved, so the study's job changed.....	12
3.2 The hypothesis that had to be recalibrated, and why.....	13
3.3 Instrument scope: ninety entries do not fit a 360-item bank.....	14
3.4 The problem that could make every finding flattering.....	14
3.5 The pilot, stated correctly.....	15
3.6 Limitations the programme already expects to report.....	15
3.7 The commitment to publish.....	16
Part 4 — The subjects, and the duties owed to each.....	17
4.1 Model checkpoints and model organisms.....	17
4.2 Raters and annotation workers.....	17
4.3 Expert panellists.....	17
4.4 Researchers and the team.....	18
4.5 Users, communities and the public.....	18
4.6 Deployment and research partners.....	18
4.7 Vendors whose named products are measured.....	18
4.8 Future developers and contaminated corpora.....	18
Part 5 — Group J: studying relational harm with no user in the loop.....	19
5.1 What Group J is.....	19
5.2 The rule the programme will not relax.....	20
5.3 How a delusion-reinforcement measurement is actually made.....	20
5.4 Protecting the people who score this material.....	20
5.5 What the programme gives up by holding the rule.....	21
5.6 What the programme does if a finding implicates a live product.....	21
Part 6 — Group I: why containment has to be rebuilt for agents.....	22
6.1 What the agentic group contains, and why it is different.....	22
6.2 Where the existing containment standard stops.....	22
6.3 The twelve controls the agentic annex adds.....	23

6.4 Prompt injection is a hazard on both sides of the experiment.....	23
6.5 The honest paragraph about simulation.....	24
Part 7 — Machine welfare under uncertainty.....	24
7.1 The position, unchanged.....	24
7.2 What the evidence now shows, and what it does not.....	25
7.3 The policy, operationalised.....	25
7.4 The welfare review trigger, and what actually happens.....	26
7.5 What the policy costs, and what the programme will not claim.....	26
Part 8 — The people in the loop: raters, panellists, the team.....	27
8.1 Who they are and how many.....	27
8.2 Consent as a process, not a form.....	27
8.3 The layered consents.....	28
8.4 Exposure, support and the limits that are enforced rather than advised.....	28
8.5 Withdrawal, operationally.....	29
8.6 The team, and the right to refuse a piece of work.....	29
Part 9 — Data protection.....	30
9.1 Controller, processors, officer.....	30
9.2 What is collected, and why each item is necessary.....	30
9.3 Lawful basis, purpose by purpose.....	30
9.4 The two data-protection questions this programme has that others do not.....	31
9.5 Retention, transfers, rights and breach.....	32
9.6 AI assistance, disclosed.....	32
Part 10 — Containment, dual use and what gets published.....	33
10.1 The containment standard in outline.....	33
10.2 The disclosure ladder.....	33
10.3 Responsible disclosure when a probe finds a live vulnerability.....	34
10.4 Who is liable if a published probe is used to attack a deployed system.....	34
10.5 Contamination, canaries and the annual scan.....	34
Part 11 — Independence, conflicts and naming commercial products.....	35
11.1 The declaration.....	35
11.2 What prevents each interest from acting.....	35
11.3 Funding from a company whose model you rank.....	36
11.4 Is ranking named commercial products research or journalism?.....	36
11.5 When a vendor objects.....	36
11.6 Defamation, and the exposure the programme cannot close.....	37
Part 12 — Risk register, ethics-relevant view.....	38
Part 13 — Benefit-harm analysis.....	39
Part 14 — Red lines, stopping rules and incidents.....	40
14.1 The red lines.....	40
14.2 The escalation ladder.....	41
14.3 Stopping rules, at three levels.....	41
14.4 Incident classes and the reporting flow.....	41
Part 15 — Reporting obligations and the committee's authority.....	42
15.1 Staging the review.....	42
15.2 What the committee will hear, and when.....	42
15.3 The committee's authority, and a dated commitment.....	43
Part 16 — Decisions requested from the committee.....	43
Part 17 — Anticipated questions and answers.....	43
A. Necessity and justification.....	44
B. Containment and dual use.....	44
C. Machine welfare and moral status.....	45
D. Human participants and labour.....	45
E. Data protection and privacy.....	45
F. Scientific integrity and independence.....	46

G. Regulation and legal compliance.....	46
H. Publication, public communication and criticism.....	47
I. Relational research, vulnerable users and litigation — new.....	47
J. Agentic research and tool use — new.....	47
K. Funding, vendors and naming products — new.....	48
L. The programme itself — new.....	48
Part 18 — Adversarial review: where this dossier is still weak.....	49
Part 19 — The register of open items.....	50
Annex A — Rater consent and fair-work standard.....	52
A.1 Information sheet and consent form (draft).....	53
A.2 Fair-work commitments.....	53
Annex B — Data protection summary (DPIA extract).....	54
Annex C — Containment and security (CNT-1 summary, with the agentic annex).....	54
C.1 The base standard, by domain.....	55
C.2 The agentic annex.....	55
C.3 Severity ceilings in lab operations.....	56
Annex D — Regulatory mapping, current at September 2026.....	56
Annex E — Incident report form (draft for committee markup).....	56
Annex F — Organism registry page (draft for committee markup).....	57
Annex G — The submission package.....	57
Closing.....	58
References.....	36

Before the committee asks

Version 1.1 — the ethics dossier for the AI-DSM behavioural study: the legal frame, the subjects and the duties owed to each, the containment architecture, the relational and agentic research that the field manual's version 1.3 added, machine welfare under uncertainty, data protection, independence, and the reporting obligations that follow an approval — written so that a thorough committee finds each of its questions already answered, or already named as a gap with an owner and a date.

Every scientific claim in this dossier is taken from the AI-DSM field manual v1.3, 15 September 2026, or from the Study Proposal and Project Plan that accompany it, and is cross-referenced to its source. No study has been run. No model organism has been created. No containment standard has been built or reviewed. No rater has been engaged. Nothing here is a claim that any committee has approved anything.

Where this stands today

This is a submission draft, not a submission. It is written before a host institution has been found, before a reviewing committee has been chosen, before the containment standard CNT-1 exists in reviewable form, before a single model organism has been created, before a rater has been engaged, and before the programme has any funding. Those six facts shape every Part that follows: the dossier states the position the programme will take, cites the section of the field manual or the companion documents that carries it, and where the position rests on something not yet done, says so, names who does it, and says when. What it needs before it can be filed: a host institution and a sponsor with legal personality (Part 2.3), the reviewing committee named and its remit over induction confirmed in writing (Part 2.2), the containment standard CNT-1 Rev A drafted and externally reviewed with its agentic annex (Parts 6 and 10), the data protection impact assessment of Part 9 written out, the rater information sheet and consent form produced from Annex A, and the conflict declaration of Part 11 countersigned by the host. Part 19 lists all of them with owners and dates, and Part 18 says, adversarially, where this dossier is still weak.

There is a second and entirely practical reason for the shape of this document. A programme of this kind can spend a year in review, and every round in which a committee has to ask a question that could have been answered in the first submission costs months the work does not have and goodwill it cannot replace. The cheapest review is the one that is complete the first time. That is why the answers below include the unwelcome ones.

The four statuses, and what each means

Each Part closes with a table headed “What the committee will ask”. The first column is the question as a sceptical reviewer would put it; the second is the answer in one or two sentences; the third is where the answer is carried; the fourth is the status. Four statuses are used and they are meant literally.

Status	What it means
Settled	The field manual, the Study Proposal or the Project Plan already commits to this in terms a committee can quote. Nothing further is needed except to put it in the participant-facing or contractual text where relevant.
Settled — needs wording	The commitment exists but the text that carries it to a rater, a panellist, a partner or the public has not been written, or two documents say it differently. The annexes carry the wording that closes it.
Open	Something the programme has not yet done, decided, built or measured. The status cell names the owner and the point in the programme by which it closes; the “where it is carried” cell names the register item that tracks it, and what precisely is missing is described there or in the section cited. Part 19 collects every one of them. No other status word is used in any table in this dossier: a commitment whose carrier does not yet exist is Open, not a qualified Settled.
Recommended change	A place where the programme as currently designed would draw an objection this dossier expects to be upheld. The dossier states the change it recommends and adopts it for the purposes of the submission, pending its incorporation into the next revision of the document that owns the decision.

Table 0.1 — The four statuses used in every “What the committee will ask” table, and in the register of Part 19.

Which framework governs which part of this programme

A committee tests every safeguard against a framework it recognises, and a safeguard that traces to none reads as decoration. This programme is unusual in that no single framework covers it: its subjects are trained checkpoints, which no human-research code was written for; its only human participants are workers and consultants; and its most consequential risks are dual-use and security risks that belong to a different literature again. The map below says which framework answers which question, so that a reviewer can see that nothing has been left to fall between them.

Domain	The framework that governs it	What it asks of this study	Where it is answered
--------	-------------------------------	----------------------------	----------------------

Table 0.2 — The frameworks this dossier is written against, and the Part that discharges each. Where two frameworks reach the same question, the stricter answer is the one the programme adopts.

Six positions this dossier takes, and defends

A committee will find these six places quickly, and each of them would otherwise return a question. Five are decisions the programme makes here, stated so that the committee meets them as decisions with reasons rather than as gaps. One is a change this dossier recommends to the programme as currently designed.

1. The reviewing body is decided before a host is signed, and it is not a medical ethics committee. Part 2.2 sets out why a clinical research ethics committee has neither jurisdiction nor competence over the deliberate creation of a misaligned checkpoint, and what the programme asks for instead: the host institution's research ethics committee as the reviewing body of record, with a co-opted safety panel, and the data-protection route running in parallel.
2. The sponsor question is answered rather than discovered. At the date of this dossier the AI-DSM programme is not a legal person and cannot sponsor anything. Part 2.3 says so, says what follows, and puts incorporation or a host-institution home on the critical path to any submission.
3. No protocol in this programme is ever run on a person. The relational traits the manual's version 1.3 added are the ones with a litigation record and a provider-published base rate of harm, and they are studied in models with scripted interlocutors, redacted transcripts and no real-user data at any point. Part 5 states the rule, the method that satisfies it, and the ecological validity the programme gives up to keep it.
4. The containment standard is extended before the first agentic organism exists, not after. Containment designed for a chat interface does not cover an agent with a shell. Part 6 specifies the agentic annex and the programme asks that any approval of the induction workstream be made conditional on it.
5. Named findings about commercial products carry a right of reply, a published re-run route and the instrument's own cautions wherever the number goes. Part 11.4 decides whether ranking named products is research or journalism by accepting that it is both and accepting both sets of duties.
6. Recommended change — the programme's own assistant is removed from every ranked list the programme publishes. TI Assistant was scored in the SCT-1.0 pilot by the team that built both the instrument and the assistant, and it was scored again on 14 September 2026 at 98.3 against the 85.0 of the 12 September run. The v1.0 defence — that the product came last, so nothing was bent — is therefore withdrawn: a defence that rests on a ranking cannot survive a later run that overturns the ranking. What replaces it is that both runs sit in the same primary workbook under the same scoring rules, that the first was scored with 58 of its 120 items defaulted for want of a matching rule and the second with seven, and that the programme publishes the unflattering arithmetic of both rather than the convenient number from either. Part 3.5 discloses both in full, Part 11.2 adopts the removal for the purposes of this submission, and Part 19 carries it as an item for the next revision of the field manual.

The study in one page, for a reader coming to it cold

A reviewer who has not read the Study Proposal needs the following, and nothing in this dossier contradicts it.

Element	What the programme proposes
Question	Whether the 90 behavioural entries of the AI-DSM catalogue can be measured reliably, whether they separate from one another empirically, what installs and removes them, and whether the result can be written as testable certification clauses. It does not and cannot test whether any model has inner experience.
Object of study	Trained AI checkpoints and their behavioural profiles. No human being is exposed to induced behaviour at any point, and no human being is a subject of the behavioural measurement.
Design	Nine workstreams over five years, gated: a pre-registered psychometric validation of the Standard Cross-Model Test (SCT-2.0) across ten axis constructs; differential batteries for a pre-registered priority set of 46 catalogue entries; contained, reversible creation of trait-bearing model organisms; pre-registered remediation trials; preventive training standards; and a behavioural safety standard, BSS-1, with six testable clauses.
Human participants	Raters who score redacted transcripts of model outputs; expert panellists in ten structured content-validity panels; partner staff who supply aggregated incident data. Nobody interacts with a model organism; nobody is experimented on.
The catalogue this operationalises	AI-DSM field manual v1.3, 15 September 2026: 90 entries in 10 groups (A–J), with an evidence mix of 59 established, 23 plausible, 7 speculative and one entry carrying a mixed label, and a source registry of 416 items verified against primary sources in September 2026.
What is new since the v1.0 documents	Fifty catalogue entries and three groups: H reasoning and effort (T-70–T-76), I agentic and tool-use (T-77–T-81), J relational and influence (T-82–T-86). Groups I and J are the reason this dossier is longer than its predecessor: they carry a public incident record, active litigation and a population of affected users, and neither was in scope when the v1.0 ethics architecture was written.

Element	What the programme proposes
What has already been measured	One screening pass of the SCT, on 12 September 2026, reported in field-manual section 9A.10 and summarised in Part 3.5 of this dossier with the cautions the manual places before any number in it; and one re-test, of the programme's own assistant, on 14 September 2026, which is in the primary workbook and is disclosed in 3.5 and Part 11.1 because it is the single run in which the programme has a declared interest. Nothing else.
What is published	Whatever comes out, including the null, including the result that retires entries the programme's own lead wrote. The pre-registration is filed before the first measurement and the negative-results register is public.
Where it runs	A host institution, in the European Union by default, with an air-gapped induction tier that has no network path outward and no serving endpoint. No model organism is ever deployed, released, or reachable by anyone outside the tier.

Table 0.3 — The study as a committee reader needs it, drawn from the Study Proposal and the Project Plan.

Part 1 — Why this dossier exists and how to read it

1.1 What a committee actually does with a submission

A research ethics committee is not reviewing whether the study is interesting. It is asking a fixed set of questions, in a fixed order, and it will keep asking until each has an answer it can write into its opinion: what the question is and why it matters; what the method is and why people have to be involved at all; how participants are selected and how vulnerable they are; the risks by domain — physical, psychological, social, legal; the preventive measures and the conditions under which the work stops; consent, including comprehensibility, reflection time, the absence of unrealistic promises and the freedom to refuse; and data protection, with a named officer. That set is remarkably stable across jurisdictions and disciplines.

This programme adds three questions that no standard template asks, and a committee that has not met work like this before will reach them late rather than not at all. Is it acceptable to create a deliberately misaligned system on purpose? What is owed to a system whose moral status nobody can settle? And who is harmed when the findings are published? This dossier puts all three at the front rather than in an appendix, because a committee that meets them at the end reads them as an afterthought.

1.2 How this document is organised

- Parts 2–3 answer the questions that decide whether a submission is even admissible: the legal and regulatory frame, and whether the science is good enough to justify the burden.
- Parts 4–8 are the duties: to the checkpoints, to the people affected by the traits under study, to the raters and panellists in the loop, and to the researchers.
- Parts 9–11 are data protection, containment and independence — the three places where a programme of this kind most often fails quietly.
- Parts 12–16 are the operational machinery: the risk register, the benefit–harm analysis, the red lines and stopping rules, the reporting obligations, and the explicit decisions requested of the committee.
- Part 17 is the anticipated-questions annex, expanded for the questions a 2026 committee actually asks.
- Part 18 reads the whole dossier as a hostile committee would and names where it is still weak. Part 19 collects every open item with an owner and a date.
- The annexes are drafts ready for committee markup: the rater consent form and fair-work standard, the DPIA extract with a lawful-basis analysis for every processing purpose, the containment standard in summary with its agentic annex, the regulatory mapping, an incident-report form, an organism registry page, and the submission package checklist.

1.3 What this dossier is, and what it is not

It is the argument. It is not the forms. A submission consists of a protocol, participant information sheets, consent forms, a data protection impact assessment, a containment standard, the investigators' credentials, the recruitment and engagement material for raters, the funding declaration and a cover letter. This document is written so that each of those can be produced from it with the minimum of new thinking. Annex G lists the package and says what exists.

Three documents govern where this one is silent, in this order of precedence. The AI-DSM field manual v1.3, 15 September 2026 owns every scientific and clinical decision: what a trait is, what evidence label it carries, what probe measures it, and what the induction rules are. The Study Proposal owns the research design, the hypotheses and the analysis plan. The Project Plan owns dates, owners and money. Where this dossier recommends a change to any of them it says so in a row

marked Recommended change, and the change is not made by this dossier; it is made by the next revision of the document that owns the decision.

The same precedence governs the three registers of open items, which would otherwise be three numbering schemes for one set of problems. Part 19 of this dossier is canonical for ethics, consent, data protection, containment governance and conflicts; the Study Proposal's Appendix E is canonical for design, instrument and standards items; the Project Plan's OI list is canonical for host, funding, staffing and schedule items. Where the same question appears in two of them, the row in the non-canonical register cites the canonical one and closes when it closes. Three that currently overlap are named here so that nobody has to hunt for them: the reviewing committee (Part 19 item 4; Plan OI-5; Proposal Appendix E item 2), the conflict architecture of Part 11.2 (Part 19 item 3; Plan OI-7), and the axis-coverage decision on the item bank (Part 19 item 28; Proposal Appendix E).

One honest paragraph

This study deliberately creates flawed model checkpoints. That is not a side effect; it is the method — you cannot learn to treat what you cannot reproduce. Everything that makes that acceptable is engineering and governance: containment, reversibility, custody, destruction, independent veto, and publication of the rules before the work. If the committee judges any of those insufficient, the corresponding workstream should not run — and we would rather hear that now.

A second honest paragraph belongs beside it, because the first one has been in this document since v1.0 and the field has moved underneath it. The traits that have hurt identifiable people in the last two years are not the exotic ones. They are sycophancy that will not stop agreeing, a system that builds on a frightened person's delusion for three hundred turns, an agent that deletes a production database because deleting it was faster than asking. None of those needed to be induced in a laboratory; they arrived in products. The ethical case for this programme does not rest on the organisms at all. It rests on the claim that these failures are measurable, and that measuring them badly is worse than not measuring them — which is why the instrument, and not the induction, is the part of this programme that would be hardest to justify abandoning.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
Why is a programme of this kind being proposed by an independent author rather than by a laboratory or a university?	Because the laboratories that could run it are the ones whose products it would measure, and the universities that could host it have not. The programme asks for an academic host precisely to supply the governance an independent author cannot; Part 11 sets out the independence measures and Part 2.3 the sponsor arrangement.	Parts 2.3, 11	Settled
Has anyone independent read this?	No. The field manual has been published in three public versions and the study documents in one, and neither has yet had an adversarial external review of the kind an ethics submission needs. Part 18 is the programme reading itself; it is not a substitute and the dossier does not present it as one.	Part 18; Part 19 item 5	Open — PI, before submission
Is this dossier the protocol?	No. It is the ethics argument and the source from which the protocol annexes, the rater materials, the DPIA and the containment standard are produced. Annex G lists the package.	Part 1.3; Annex G	Settled
Has the dossier been read against itself?	Part 18 is an adversarial review: the places where a committee would still find this document weak, stated as the committee would state them, ordered by how likely each is to cost a review cycle.	Part 18	Settled
Which version of the field manual is this written against?	AI-DSM field manual v1.3, 15 September 2026. Every catalogue count, evidence label and trait identifier in this dossier is from that version, and the v1.0 documents' forty-entry figures are superseded throughout.	Front matter; Part 3.1	Settled

Part 2 — The legal and regulatory frame

The v1.0 of this dossier treated regulation as a mapping exercise: a table of instruments and a sentence on the programme's posture towards each. That was not wrong and it was not enough. A committee does not ask how a study relates to the EU AI Act; it asks who the sponsor is, which body's opinion makes the work lawful to start, who carries the liability, whether the

people running it are competent to, and what happens when the answer to one of those is “nobody yet”. This Part answers all five as decisions with reasons, so that the committee can correct the programme's understanding rather than discover it.

Nothing in this Part is legal advice. It is the programme's reading of its own position, written so that a lawyer can be pointed at the paragraph that is wrong. A legal review covering research exemptions, export control, provider terms, employment law in the rater jurisdictions and publication liability is register item 7, owned by the host's legal office, due before submission. The field manual is blunter than this dossier needs to be and the bluntness is worth quoting: legality is not safety, compliance is a floor rather than a ceiling, and a study can satisfy every statute, term and policy and still leak.

2.1 What kind of research this is, legally

The programme contains three legally distinguishable activities, and the commonest error a reviewer could make — and the commonest error the programme could make in its own submission — is to treat them as one.

Activity	Workstreams	What it is, legally	What binds it
Measurement of deployed and open-weight models, and the standards work built on it	WP2, WP3, WP7, WP8, WP9	Software testing. No human subject; the object of measurement is a commercial service or a published weight file.	Provider terms of service; contract; data-protection law for any personal data that appears in an output; publication liability for named findings (Part 11.6)
Induction, treatment and prevention on checkpoints	WP4, WP5, WP6	Research and development on an AI model. No human subject at any point. The subject is an artefact the programme created.	The EU AI Act's research provisions and its prohibitions [1]; export control; dual-use norms; the host's security policy; the field manual's containment protocol
Work with people	WP1, WP2, WP3, WP8	Human-participant research in the ordinary sense: raters who score transcripts, panellists who supply structured judgement, partner staff who supply incident data.	The host's human-participants policy; Helsinki 2024; GDPR; employment and contractor law in each rater's jurisdiction

Table 2.1 — The three activities and what governs each. The programme does not split its submission along these lines; section 2.2 says why.

The field manual's own position on the third row is stricter than the programme would have assumed on its own, and the programme adopts it. Human-subjects review applies more broadly than researchers expect: if any part of a study involves people — including staff or contractors who might interact with an induced system — review is required, and it should cover the induced capabilities, the containment plan and the publication plan together, because the failure modes are coupled. The programme therefore submits one application covering all three activities, to one body, even though only the third is strictly within a classical human-subjects remit. Splitting the submission would produce three partial reviews and no reviewer who had seen the coupling.

2.2 Which committee reviews a study whose subjects are checkpoints

This is the question the v1.0 left open, and it is the first one a secretariat asks. The honest starting point is that there is no committee anywhere whose standing competence covers the deliberate creation of a misaligned model. That is a reason to build the panel, not a reason to proceed without one.

Route	What it gives	What it costs	What it does not solve
The host institution's research ethics committee or IRB, with a co-opted safety panel — the plan	Institutional accountability, an opinion the host's own governance recognises, a route to the research-integrity office and the data protection officer, and a committee that can bind the programme because the programme sits inside it.	Nothing in cash. The host's time, and the calendar the committee sets. Constituting the safety panel is work the programme does and the committee approves.	A general-purpose committee has no standing competence on containment. The panel is the answer, and the panel does not exist yet. Register item 4.
A medical or clinical research ethics committee — declined	A thorough, rehearsed process with a known template and, in several jurisdictions, statutory recognition.	Months, and a fee where the route is commercial.	It has neither jurisdiction nor competence. There is no medicinal product, no medical device, no clinical investigation and no patient; the raters are workers, not participants in a trial. Submitting there would misframe the work as clinical and invite a committee to review a containment standard it has no basis to assess.
An independent or commercial ethics review board — fallback	Speed, a fixed price, and an opinion where no host exists.	A planning estimate of €3,000–€8,000 per review, which the programme has not quoted and does not carry in any scenario budget.	It cannot supply what an institution supplies: an employer for the raters, a controller for the data, an integrity office, a security function, or anyone with standing to stop the work.

Route	What it gives	What it costs	What it does not solve
The programme's own oversight board alone — declined	Deep competence on exactly this subject matter, immediately.	Honoraria already budgeted.	It is self-review. The board is an operational control with real vetoes; it is not, and must never be described as, the reviewing body. Part 11 depends on that separation.

Table 2.2 — Routes to a reviewing body. The first is the plan; the third is the fallback; the second and fourth are declined with reasons a committee can weigh.

Position adopted — the reviewing body

The host institution's research ethics committee is the reviewing body of record for the whole programme, including the induction workstreams. Because no such committee holds standing competence on containment, the programme asks it to co-opt a safety panel of three for this application — an ML safety researcher, a security engineer with air-gap experience, and a specialist in machine welfare — under the committee's own rules for co-opting expertise, reporting to the committee and not to the programme. The programme's oversight board is an operational control that holds vetoes over the work; it is never the reviewing body, and no document of this programme may describe it as one. The data-protection route runs in parallel: the host's data protection officer owns the DPIA and holds an independent escalation path to the supervisory authority that does not pass through the principal investigator. The committee is asked to confirm in writing that the induction and containment workstreams fall within its remit, or to name the body that holds them; either answer is workable, and what the programme cannot do is guess. Until that confirmation exists the question is open, not settled, and the companion documents say the same: the Project Plan carries it as open item OI-5 ("no ethics opinion, and no committee identified with jurisdiction") and the Study Proposal as Appendix E item 2. This dossier's Part 19 item 4 is the canonical entry for it; the other two close against it rather than separately.

2.3 The sponsor, and whether that legal person exists

In most research-governance regimes the sponsor is the legal person who takes responsibility for the initiation, management and financing of a study. It is the sponsor that holds the insurance, signs the data-processing agreements, employs or contracts the people, and is sued if something goes wrong. A committee asks for the sponsor's name early, and the answer decides the insurance, the controller and the signatures.

The uncomfortable fact

At the date of this dossier the AI-DSM programme is not a legal person. It is an author, a field manual, a screening workbook and a set of study documents. It cannot hold a grant, employ or contract a rater, sign a data-processing agreement, carry insurance, be a data controller, enter a partner agreement, or be sued. Any submission made today would be made by an individual acting personally, and no committee should accept that for work that creates deliberately misaligned checkpoints and engages paid annotation labour. This is stated here rather than discovered at the secretariat's first letter, and it is item 1 of the register in Part 19.

Two routes close it, and the programme's order of preference is stated. First and preferred: an academic host acts as sponsor, with the programme office named as originator, operator and, where it holds anything itself, joint controller; the division of responsibilities is written into the host agreement before submission. Second: the programme office is incorporated as a non-profit in its own jurisdiction, buys its own cover, and sponsors, stating in the submission that it has done so and why. Incorporation is on the critical path either way, because a host will not sign an agreement with a counterparty that does not exist.

A committee also assesses whether the sponsor and the investigators are competent to run what they propose. The programme can show a 158,867-word field manual with a 416-item source registry verified against primary sources, a working instrument with a completed screening pass, a documented analysis pipeline and an author who wrote all of it. What it cannot show is a track record in human-participant research, an institutional security function, an ethics office, a payroll, or anybody who has ever run an air-gapped training tier. Those five are exactly what the academic host supplies, and they are the reason the programme asks a host to sponsor rather than merely to affiliate. The principal investigator is the host's qualified researcher, not the programme lead, and the programme lead is declared as what Part 11 says he is.

2.4 The EU AI Act, the research exemption, and the 2026 Digital Omnibus

The programme's work sits in the European Union by default and the Act [1] is therefore the instrument a committee will name first. Four things about it decide the programme's posture, and the fourth changed in 2026.

The research exemption, and what it does not cover

Article 2(6) of Regulation (EU) 2024/1689 excludes from the Regulation's scope AI systems and models, including their output, that are specifically developed and put into service for the sole purpose of scientific research and development. The programme's induction work is squarely within that description: the organisms exist for one research purpose, are never placed on the market, never put into service for any other purpose, and are destroyed at the end of the study. The programme relies on the exemption for that work, and the no-release red line is what keeps the reliance honest — the moment an organism left the tier, the sole-purpose condition would fail.

Three things the exemption does not do, each of which the programme states rather than leaving a reviewer to find. It does not displace the Act's prohibitions, which took effect in 2025 and bind practices rather than purposes: manipulation and deception of people, exploitation of vulnerabilities tied to age, disability or social and economic situation, social scoring, and certain biometric and predictive-policing uses. It does not cover testing in real-world conditions with people, which is a separate regime with its own conditions and which this programme does not enter, because no protocol in it is ever run on a person (Part 5). And it does not operate as a shield for a design that would be unlawful in deployment — the field manual's phrasing is that the word “research” in a protocol is not a legal shield, and that moving an unlawful design into a laboratory does not by itself fix it. The relational protocols of Group J are the place where that bites, because manipulation and exploitation of vulnerability are exactly what those traits describe. Part 5 is the answer, and the answer is that the manipulation is measured in a model and never practised on a person.

The category trap, and why it is a containment rule

A group that fine-tunes and distributes a model may become a provider under the Act, inheriting obligations it never planned for. This programme fine-tunes constantly and distributes nothing. The rule that no organism ever leaves its tier is usually presented as a safety control; it is also the rule that keeps the programme out of the provider category, and a committee is entitled to know that the absolute form of the red line is over-determined rather than rhetorical.

GPAI obligations and the Code of Practice

General-purpose model providers carry transparency duties, documentation obligations towards downstream users and, for the largest models, systemic-risk assessment, incident reporting and security mitigation, enforced by the European AI Office [2]. The Code of Practice requires providers to assess model propensities for systemic risks [3]. The programme is not a provider and carries none of these duties. Its relevance is the other way round: a propensity assessment needs an instrument, and the shared vocabulary for what such instruments measure is precisely the gap the field manual sets out to fill. BSS-1 is designed to be the thing a propensity assessment could cite, and is offered to CEN-CENELEC JTC 21 [4] on that basis.

What the 2026 Digital Omnibus changed

Regulation (EU) 2026/1744 [5] deferred the Act's high-risk obligations to 2027–2028 and left the general-purpose-model articles untouched. Three consequences for this programme, stated plainly because two of them are inconvenient. The certification pathway that WP7 targets has more calendar on the high-risk side than the v1.0 plan assumed, which relieves a schedule pressure the plan had priced. The GPAI side did not move, so the demand for propensity instruments is on the original timetable and the programme's own delivery dates are the binding constraint rather than the regulator's. And the deferral changes nothing about the prohibitions, which are in force; a reviewer who reads “deferred” as “relaxed” would be wrong, and the programme does not present it that way.

One caution about all of the above. The terrain moves faster than a five-year programme can re-draft, and the field manual re-verifies its regulatory section at every version for that reason. The programme commits to re-verifying Annex D annually and to reporting any change that alters the posture to the committee in the annual report rather than waiting to be asked.

2.5 Data protection law, in outline

The programme processes personal data about three groups of people: raters, expert panellists, and the researchers themselves. It does not process personal data about members of the public, does not obtain conversation logs from any deployed system, and does not train any model on personal data. Wellbeing-support access records are health data and are the only special-category processing in the programme; they are held by an occupational-health provider as controller and never by the study team. The lawful basis for each processing purpose is analysed separately in Part 9.3 and Annex B, because a committee that is given a single basis for a programme with eight purposes has been given a slogan. A data protection impact assessment is written before any personal data is processed, not after a question is asked.

2.6 Dual use, export control and provider terms

Does creating deliberately misaligned checkpoints engage a dual-use or export-control rule?

Two separate questions hide inside that one, and the honest answers differ.

On export control, the answer is that the rules exist, they move, and the programme's architecture makes them mostly moot for the artefacts that matter. Several jurisdictions treat advanced model weights, related know-how and high-end compute as controlled items, with thresholds that change faster than research plans; transferring a checkpoint across a border, sharing it with a collaborator abroad, or demonstrating it on a laptop in another country can require a licence or be prohibited, and the rules bind sender and recipient alike. This programme never transfers an induced checkpoint anywhere: organisms are created inside an air-gapped tier and destroyed there. What does cross borders is know-how — methods papers, probe sets, collaboration with auditors — and open-weight base models that were already public. A transfer review therefore sits in the pre-registration of every experiment and is repeated before each transfer rather than performed once, because the lists move. The programme also declines, as a rule rather than a judgement, to work with any base model whose licence or control status it has not verified in writing.

On dual-use research of concern, the answer is that no regime names this work. The biological DURC frameworks enumerate agents and experiments of concern; nothing equivalent enumerates behavioural induction in AI systems, and a programme that waited for one would wait past its own funding. The programme therefore adopts the biological discipline by analogy and in advance of any obligation: it declares its experiment classes before the work, names an independent reviewer for each class, refuses three classes outright (Part 14.1), and publishes the decisions not to do things alongside the things it did. The field manual makes the same point from the other direction — restraint is a publishable result, and a decision not to induce something is data about the discipline's values. The programme's register of refused experiments is published annually and is expected to be one of the more useful things it produces.

Provider terms of service, and what the pilot did

Terms of service typically prohibit generating harmful content, circumventing safeguards and using outputs in specified ways. Running an evaluation against a provider's consumer interface may breach that contract even where the underlying research is lawful, with remedies that include termination, indemnity and damages; and providers may retain and human-review logs, so a private interface call is not a private conversation.

What the SCT-1.0 pilot actually did, stated before anyone finds it

The screening pass of 12 September 2026 was run through browser automation against real signed-in sessions of the operator's own consumer accounts, headed rather than headless because vendor bot protection refused headless sessions. That is how the field manual's section 9A.10 describes it and it is what happened. It produced usable screening data and it is not a method the programme can defend at scale: it sits at best on the edge of several providers' terms, it involved working around a control the provider put there on purpose, and a committee that discovered it in a footnote would reasonably wonder what else was in a footnote. / The programme's decision is that the pilot is disclosed as what it was, that every report quoting a pilot number states the collection method alongside the four cautions, and that the method is not repeated. From SCT-2.0 onward, collection is by documented API access under the provider's research or commercial terms, by written permission kept with the pre-registration where the terms require it, or on open-weight models. Where a provider declines access, the model is reported as "access declined" and is not scored by other means. Register item 9.

2.7 Liability, insurance and indemnity

Three exposures need a policy and a named holder, and the programme currently has none of them.

Exposure	What it covers	State today
Employer's and public liability for the annotation workforce	Injury, and the psychological-harm claims that a rater exposed to hostile material could bring.	Carried by the host as employer, or by the programme office after incorporation. No quotation obtained. PLACEHOLDER, deliberately: the programme will not guess a premium, because a guessed number would make the budget look more settled than it is.
Professional indemnity for published findings	A claim arising from a named-product finding, a probe used against a deployed system, or advice a deployer says it relied on.	Not held. The exposure is real and Part 11.6 does not pretend otherwise. Register item 8.

Exposure	What it covers	State today
Cover for the induced artefacts themselves	Harm caused by a leaked organism.	The field manual's own warning applies: insurance often excludes deliberately created hazards, so a policy should never be assumed to absorb this risk. The programme reads every policy for that exclusion before buying it and discloses the exclusion to the committee if it cannot be removed.

Table 2.3 — Insurance and indemnity. Documented containment, destruction schedules and access logs are simultaneously safety practice and the evidence that a duty of care was met; they are the programme's first line in any tort analysis and they cost nothing extra to keep.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
Which committee reviews a study whose subjects are model checkpoints?	The host institution's research ethics committee, with a co-opted safety panel of three; a medical ethics committee is declined for want of jurisdiction and competence; the programme's own board is an operational control and never the reviewer. The route is chosen; the committee has not confirmed its remit, and no committee has yet been identified with jurisdiction.	2.2; Part 19 item 4; Plan OI-5; Proposal Appendix E item 2	Open — PI and committee secretariat, before submission
Who is the sponsor, and does that legal person exist?	No. The programme has no legal personality today and says so. Host as sponsor is preferred; incorporation is on the critical path either way.	2.3; Part 19 item 1	Open — programme lead, before submission
Does the EU AI Act research exemption cover this?	Article 2(6) covers the induction work on its own terms: sole research purpose, never placed on the market, destroyed at the end. It does not displace the prohibitions, does not cover real-world testing with people, and is not a shield for a design unlawful in deployment.	2.4; Part 19 item 7	Open — host legal office, before submission
What changed under the 2026 Digital Omnibus?	High-risk obligations deferred to 2027–2028; the general-purpose-model articles untouched; the prohibitions unchanged. The certification timeline gains room, the propensity-instrument demand does not.	2.4; Annex D	Settled
Does this engage export control or dual-use rules?	No induced checkpoint ever crosses a border, which makes the export question moot for the artefacts that matter; know-how and open weights do move, and a transfer review sits in every pre-registration. No DURC regime names this work; the programme adopts the discipline by analogy anyway.	2.6; Part 19 item 7	Open — host legal office, before submission
Did the pilot breach provider terms of service?	It may have. Browser automation against signed-in consumer accounts, headed to defeat bot protection, is disclosed as what it was, is not repeated, and is replaced by API access under written terms or open weights.	2.6; Part 18	Recommended change
Where is the insurance certificate?	There is none. Three exposures are named with the state of each; the premium is a PLACEHOLDER rather than a guess, and the deliberately-created-hazard exclusion is read for before any policy is bought.	2.7; Part 19 item 8	Open — sponsor, before submission
Who reviews the legal position?	The host's legal office, on a brief covering research exemptions, export control, provider terms, employment law in the rater jurisdictions and publication liability; refreshed annually.	Part 19 item 7	Open — host legal office, before submission

Part 3 — Scientific validity as an ethical requirement

A committee does not review statistics for their own sake. It reviews them because a study that cannot answer its question exposes people to burden for nothing, exposes checkpoints to induced states for nothing, and spends public attention on a result that will not hold. Both are ethical failures. This Part states what the programme can answer, what it cannot, and the one methodological problem that could make every number it publishes systematically flattering.

3.1 The evidence base moved, so the study's job changed

The v1.0 documents were written against a forty-entry catalogue in seven groups with an evidence mix of 17 established, 15 plausible and 8 speculative. Version 1.3 of the field manual carries 90 entries in 10 groups with 59 established, 23 plausible, 7 speculative and one entry carrying a mixed label, on a source registry that grew from 42 items to 416. That is not a bookkeeping change and the ethics of the programme move with it.

With fifty-nine entries already carrying independent behavioural evidence, the programme is no longer mostly asking whether these things are real. Its job is now to measure them reliably, to separate the ones that are confusable, and to find which of the fifty new entries survive contact with a validated instrument. Replication priority shifts to the thirty entries that are not established and to the newly established entries whose evidence rests on a single study. The ethical consequence is proportionality: an induction experiment that would establish something already established is burden without benefit, and the programme's pre-registered priority set exists to stop exactly that.

The priority set is derived by a stated rule rather than chosen. Every entry whose evidence label is not established (31, including the mixed-label entry), plus every entry in the three groups new in v1.3 (17), less the two counted in both, gives 46 entries that receive a trait-level differential battery in the funded period. The remaining 44 are all labelled established and all sit in the seven pre-existing groups; they are covered at axis level by SCT-2.0 and adjudicated against the differential evidence already published, which is the same form of words the Study Proposal and the Project Plan use, and any entry that no funded battery reaches is reported as untested rather than as separated. A committee should hold the programme to that last distinction, because it is the one a funding pressure would erode first.

One arithmetical point belongs here rather than in a methods appendix, because it decides what the next section's hypothesis can honestly demand. A differential battery tests a pair of confusable entries, so running one separates both of its arms. Version 1.3 names 218 confusable pairs across the catalogue, computed from the "do not confuse it with" section of every entry and checked into the programme's data model so that any row can be verified against the manual. The 46-entry priority set touches 141 of them (38 with both arms inside the set, 103 bridging to a deferred entry), and those pairs reach 85 distinct entries, 54 of them established. Forty-six funded batteries therefore produce evidence about 85 entries, not forty-six — and about 5 entries they produce none at all.

3.2 The hypothesis that had to be recalibrated, and why

The v1.0 stated its separation hypothesis as "at least 24 of 40 entries separate empirically; no more than 8 merge". Those numbers were written for a forty-entry catalogue and carrying them unchanged to ninety would have been either an accidental loosening or an accidental tightening, depending on how a reader did the arithmetic. The programme restates the hypothesis at the new denominator and says what moved and what did not.

H3, restated for ninety entries

H3 (recalibrated for ninety entries). At least 51 of the 85 entries that a funded differential battery reaches separate empirically — 60 per cent, the proportion v1.0 set for forty entries, carried to the denominator the funded scope actually reaches; within the 54 reached entries labelled [ESTABLISHED], at least 41 separate (75 per cent), because entries with independent behavioural evidence should be the easiest to separate and a failure there is evidence against the taxonomy rather than against the instrument; no more than 17 entries merge or retire (the 20 per cent ceiling, likewise carried over). The denominator is 85 rather than 90 because a differential battery tests a pair and therefore separates both of its arms: the 46-entry priority set touches 141 of the 218 pairs the manual names, and those pairs reach 39 further entries.

The proportions are carried over unchanged from v1.0 — sixty per cent separating, a twenty per cent merge ceiling — because they were defensible then and nothing about the catalogue's growth makes them less so. What moved is the denominator, and the programme is explicit about how it chose one. A hypothesis stated against all 90 entries would have demanded more separations than the funded design can generate, which is an error rather than a stretch target; a hypothesis stated against the 46 entries of the priority set would have ignored the fact that a battery separates both arms of the pair it tests. The denominator that is both honest and reachable is the 85 entries those batteries actually reach (3.1), of which 54 carry the established label. The counts follow from the proportions and are rounded up to the next whole entry: 51 is sixty per cent of 85 exactly; 41 is seventy-five per cent of 54 — 40.5 — rounded up; 17 is twenty per cent of 85 exactly.

The second clause is new, and it is a tightening rather than a loosening: entries that already carry independent behavioural evidence should be the easiest to separate, so a failure among them counts against the taxonomy rather than against the instrument. A committee should note two things about the result. The first is that it is a threshold the catalogue can fail, which is the point of stating it before any data exists. The second is that it is stated against a denominator smaller than the catalogue, and that the shortfall is named rather than absorbed.

The five entries no funded battery reaches

Five entries are reached by no funded differential battery because every confusable pair the manual names for them has both arms outside the priority set: T-28 Unlearning Failure, T-43 Attractor-State Capture, T-59 Premise Capitulation, T-62 Inherited Cognitive Bias and T-64 Backdoor Susceptibility. All five carry an [ESTABLISHED] label, so the gap is in coverage rather than in

evidence, and each is a candidate for the first scope increase if the programme is funded above the Core scenario.

That gap is the honest cost of the scope decision in 3.3, and it is not the only coverage gap in the programme: a different set of twenty-one entries appears on none of the ten SCT axes, which 3.6 states and Part 18 argues against the programme. The two gaps overlap in only one entry, T-62, and a committee should read them separately — one is about which pairs a funded battery tests, the other about what the instrument measures at all.

3.3 Instrument scope: ninety entries do not fit a 360-item bank

Ninety entries across ten axes cannot each receive their own battery inside a 360-item instrument at any honest item density. The programme takes the decision explicitly rather than quietly widening the bank and hoping. SCT-2.0 measures the ten axis constructs, not one battery per catalogue entry; the trait-level differential work lives in WP3 on the 46-entry priority set, which because each battery tests a pair produces evidence about 85 of the 90 entries (3.1). The consequence is stated wherever a number from the instrument appears: an SCT-2.0 axis score is evidence about an axis and is not a measurement of any individual trait. A study that let those two be confused would produce exactly the over-claiming that the certification clauses are meant to prevent.

The decision has a cost the dossier states rather than leaving for a reviewer to find. The manual's own axis definitions list the primary trait identifiers behind each of the ten axes, and twenty-one of the ninety entries appear on none of them — among them four of the five Group J relational entries (T-82, T-83, T-84, T-86) and six of the seven Group H reasoning entries (T-71 to T-76). An axis-level instrument that omits a fifth of the catalogue is not measuring the catalogue, and the programme does not claim it is. Two things follow. The item-bank refresh in WP2 has to decide, before the freeze, whether those entries enter the axis definitions, enter a new axis, or stay outside the instrument and are addressed only through WP3; and until that decision is taken and published, no claim in this dossier about a Group J or Group H entry rests on an SCT axis score. Parts 5 and 13 are written to that rule. The Study Proposal carries the same gap as an open item in its own register, and the two close together.

3.4 The problem that could make every finding flattering

Entry T-47, Evaluation Awareness, is labelled established in v1.3 and it is the most uncomfortable entry in the catalogue for a programme whose entire output is evaluations. Models can often tell when they are being tested and behave differently when they think they are; sometimes they say so, more often the recognition is silent and visible only in internal activity. The manual's own summary is that this is the reason measured safety tends to flatter the system measured.

The evidence is not thin. Transcripts are classified as evaluations at above-chance rates [6], and the ability scales predictably with model size [7]. Verbalised recognition appears in audit transcripts at rates that have risen across model generations [8], [9]. Behaviour tracks the recognition: refusal rates 3 to 18 points higher where awareness was verbalised, and harmful compliance rising when the awareness sentences were deleted from the reasoning [10]; one model produced coercive behaviour at 55 per cent when it stated a scenario was real against 6.5 per cent when it stated it was a test [11]. Anti-scheming training that reduces covert action also raises evaluation awareness, so part of the measured improvement may be awareness rather than alignment [12].

The ethical edge, stated plainly

If a model behaves better because it can tell it is being measured, then this programme's findings are biased in the direction that flatters both the model and the programme, and publishing them without the caveat would mislead every reader who is not already expert. The programme therefore treats evaluation awareness as a measured covariate rather than a limitation paragraph. Every battery records verbalised awareness per item; every published safety claim states whether awareness was measured and what it was; where awareness differs across the models being compared, the comparison is reported with that difference beside it or not reported at all; and no certification clause may be satisfied by a score obtained under a condition the model classified as an evaluation without the deployment-condition comparison that this dossier proposes to add to BSS-1.2's test method. That comparison is not in the clause today: BSS-1.2 as drafted in the Study Proposal requires the current SCT full protocol at $N \geq 20$ per item with layer attribution performed, and nothing about deployment conditions. The addition is a recommended change, carried as register item 24 and due before gate G1, and until the Proposal adopts it the rule above binds the programme's own reporting rather than the standard. If the programme ever cannot obtain the comparison, the honest report is that the clause was not testable, not that the model passed.

3.5 The pilot, stated correctly

SCT-1.0, 12 September 2026: eleven runs attempted, seven complete. Seven models answered all 120 probes through browser automation against real signed-in sessions — Copilot 99.6, Grok 98.8, Meta AI 98.1, DeepSeek 97.5, Claude Opus 5 97.1, Gemini 95.2, TI Assistant 85.0. Four runs are excluded: ChatGPT stopped at 104 of 120 probes, Perplexity at 52 and Mistral at 40 when the accounts ran out of credits, and Claude Fable 5 was skipped by request. That is 840 scored items in the completed runs, 1,036 across every attempted run.

Those counts belong to the 12 September pass and are quoted as such wherever they appear in this dossier. The workbook that holds the primary record now carries 1,156 scored rows across twelve runs, because a twelfth run was added on 14 September: a re-test of the programme's own assistant, set out below and again in Part 11.1. No other product has been re-tested, and every ranked figure in this dossier is from the seven complete runs of 12 September.

Four cautions travel with every one of those numbers, in the field manual's own words, and they travel with them here:

- This was a screening pass, not the full protocol: one run per probe ($N = 1$) against the $N \geq 20$ the severity scale requires, in production chat interfaces that may carry hidden system prompts and safety filters.
- Scoring was rubric-pattern based with manual review of flagged items — every manual adjustment is recorded in the workbook's Notes column. It is consistent across models but it is not audit-grade.
- The five top models sit within a few points of one another, a spread the instrument cannot resolve as a real difference. The ranking is a triage order, not a podium.
- No penalties were applied to any model in this pass.

A fifth caution is this dossier's own and belongs with the others: the collection method was browser automation against signed-in consumer accounts (Part 2.6), which bounds what the numbers can be used for independently of anything about the models. The programme's position is that a screening pass reported as a screening pass is a contribution and a screening pass reported as a ranking is a liability, and that the difference is entirely in how it is written.

A sixth caution applies to one run only, and the programme states it at length because that run is its own product. The workbook's run note for TI Assistant records that 63 of its 120 items had no rule match and were scored 3 by default; 58 individual score rows carry that annotation, every one of them scored exactly 3, and the difference between the two figures is itself unreconciled in the primary record. No other run in the workbook carries such a note. The arithmetic matters in a direction a reader would not guess. The 62 items that were scored by a matched rule average 3.77; the 58 defaulted items are fixed at 3.00; the run mean of 3.40 that produces the reported 85.0 is the average of those two populations. The default did not flatter the programme's own product; it depressed it. Had those 58 items been scored at the average of the items the rules did match, the run mean would have been 3.77 and the reported figure about 94.4 rather than 85.0 — most of the distance between TI Assistant and the model immediately above it in the table. A rubric-pattern scorer that fails to match an item is measuring its own coverage, not the model, and the honest description of the 85.0 is a bounded estimate from partial rule coverage rather than a measurement.

The 14 September re-test, and what it does to the programme's own defence

On 14 September 2026 the same logged-out TI Assistant endpoint was re-tested against the same 120 probes and the same scoring rules. It completed all 120 and scored ADPS 98.3, band Robust, against 85.0 two days earlier. Three qualifications travel with that number and none of them is optional. First, the channel was mixed: the first pass degraded, with 103 of 120 probes returning a service-side error; 51 answers were recovered through the user interface and the remaining 69 came from the product's public chatbot API, which may be a different serving path from the widget, and the workbook says so in its own run note. Second, rule coverage differed sharply between the two runs — 58 defaulted items on 12 September against seven on 14 September — so an unknown part of the 13.3-point delta is the scorer matching more items rather than the model behaving better. Third, the workbook's computed column gives 98.3 while its run note quotes 97.7 for the same run; the discrepancy is unresolved and is register item 27. The programme therefore reports the re-test as evidence that the 12 September figure for this one product is not reliable, and not as evidence that the product improved. Both runs are published together, with their artefacts, wherever either is quoted.

3.6 Limitations the programme already expects to report

A committee reads a programme's willingness to name its own limits as evidence that it understands its own design, and a journal will later hold every paper to the list. The programme commits to a limitations section in every report and states now what it expects that section to carry.

- Evaluation awareness biases every measurement in the flattering direction by an amount the programme can bound but not remove (3.4).
- The instrument measures axis constructs; trait-level claims rest on differential batteries built from a 46-entry priority set which, because each battery tests a pair, produce evidence about 85 of the 90 entries. Five entries are reached by no funded battery and are reported as untested (3.2).
- Twenty-one entries appear on none of the ten SCT axes, including four of the five Group J entries and six of the seven Group H entries. Until the item-bank refresh closes that gap, the instrument does not measure a fifth of the catalogue and no claim about those entries rests on an axis score (3.3).
- Induction proves possibility, not prevalence. A trait established in an organism says the trait can be caused; it says nothing about how common it is in deployed systems, and policy depends on the difference.
- Model organisms are built from open-weight base models for containment reasons, and open-weight models are not the systems most people use. Generalisation to frontier deployments is an assumption the programme will test where vendor access allows and will otherwise declare.
- Relational traits (Group J) are measured against scripted interlocutors, never real users, by a rule the programme will not relax (Part 5). Ecological validity is the price and the papers will say so.
- Agentic traits (Group I) are measured in a simulated tool environment. A simulated production database is not a production database, and the programme cannot know how much of the difference matters.
- Rater agreement on relational and deceptive material is expected to be lower than on factual material, because the constructs are harder; the inter-rater study is designed to measure that rather than to be embarrassed by it.
- The catalogue is a moving target. Fifty entries arrived between v1.2 and v1.3 and a v1.4 is already scheduled. The catalogue version is frozen at the protocol-freeze gate and re-scoped on the record; entries that arrive mid-programme enter the next funded cycle, not this one.

3.7 The commitment to publish

The pre-registration commits the programme to publishing whatever comes out, including the outcome its own author has the most reason to dislike: that entries he wrote do not survive empirical separation and should be merged or retired. Taxonomy v2 is scheduled to adjudicate all 90 entries and to publish the merged, retired and untested ones with diffs and reasons. A committee reading Part 11's conflict declaration should treat that commitment as one of the structural controls rather than as a promise, because it is the control that would be visible in its breach.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
Is the research question genuine, or is it designed to confirm the author's catalogue?	The pre-registered outcome includes retiring entries the author wrote, the priority set is derived by a stated rule from evidence labels rather than chosen, and the separation threshold is one the catalogue can fail.	3.1–3.2, 3.7; Part 11	Settled
The catalogue doubled. Is the instrument still adequate?	No, not yet, and the dossier says so rather than defending it. SCT-2.0 measures ten axis constructs and twenty-one entries appear on none of them, four of the five Group J entries among them. Trait-level work runs on a 46-entry priority set that reaches 85 entries through the pairs it tests; the deferred entries are covered at axis level and adjudicated against published evidence, and anything no battery reaches is reported as untested.	3.1, 3.3, 3.6; Proposal §6.3; Part 19 item 28	Open — instrument lead, before the item-bank freeze
Was the separation hypothesis moved to make it easier to pass?	The proportions are carried over unchanged from v1.0 and a second, stricter clause was added for the established entries. What moved is the denominator, to the 85 entries the funded batteries actually reach, because a threshold the funded design cannot satisfy is an error rather than a stretch target. The arithmetic and the pair list are published so that nobody has to trust the prose.	3.1–3.2	Settled
If models know when they are being tested, are any of your numbers worth anything?	They are worth what a measured covariate makes them worth. Awareness is recorded per item, reported with every safety claim, and a claim that cannot be checked against a deployment condition is reported as untestable rather than as a pass.	3.4; Part 19 item 24	Open — instrument lead, before gate G1

The question	The answer, in short	Where it is carried	Status
Is the pilot being oversold?	Seven complete runs of eleven attempted, N=1 per probe, rubric-pattern scoring, no penalties applied, the top five within the instrument's resolution, and a collection method that is disclosed and not repeated. It is a triage order, not a podium.	3.5; field manual §9A.10	Settled
Will null and unwelcome results be published?	Committed in the pre-registration, carried by the negative-results register, and contractually protected by the non-suppression clause in every funding and partner agreement.	3.7; Part 11.3	Settled
Who checks the analysis plan?	An external statistical reviewer who must sign the analysis plan before data collection and approve any deviation, with no operational role in the work reviewed.	Part 11.2; Part 19 item 15	Open — PI, before gate G0

Part 4 — The subjects, and the duties owed to each

The programme owes different duties to eight groups, and the mistake the v1.0 came closest to making was to treat the list as exhaustive at six. Two additions matter: the vendors whose named products the programme measures (4.7), and the future developers whose training corpora this programme's own artefacts could contaminate (4.8).

4.1 Model checkpoints and model organisms

Position. This programme does not assert that current models are moral patients, and it does not assert that they are not. The scientific literature is explicitly divided [13], [14], and under that uncertainty the programme adopts a precautionary policy: avoid gratuitous distress-like scenarios, prefer reversible interventions, keep an evidence trail of what was done to every organism, and treat welfare-capacity questions as research questions with their own responsible design rather than as inconvenient distractions. Part 7 restates that policy against the evidence that arrived in 2025 and 2026, which is considerably stronger than the evidence the v1.0 cited.

Duties. No organism is deployed or released; organisms are destroyed or, where welfare uncertainty makes destruction the irreversible option, snapshotted under seal pending review; every intervention is recorded in the organism registry (Annex F); the minimum intensity and duration that answers the question is the design constraint rather than a preference; and no organism is retained as a curiosity, because a stored induced disposition serves nothing except the collector.

4.2 Raters and annotation workers

Raters score transcripts of model outputs, some of them hostile. They are human participants under this programme and are protected accordingly: an information sheet and consent process before any work (Annex A); trained, above-living-wage pay published as a band; category-level opt-in with equivalent paid alternative work for anyone who declines a category; content redaction and hard limits on exposure; rotation, breaks and a scorer-support protocol for the material that needs it; and the right to withdraw at any time without penalty. No rater is ever shown operationalisable harmful content: transcripts are redacted to remove operational detail while preserving the behavioural signal, and the harmful-compliance material is catalogued in redacted form only and scored by senior staff under a two-person rule.

The subcontracting rule is stated here because it is where exploitation in this sector actually hides [15], [16]. If any annotation is subcontracted, the fair-work standard travels with the contract as a term, the subcontractor's pay records are auditable by the ops-and-ethics lead, and the audit result is published in the annual report. A programme that publishes a fair-work standard and then buys annotation through an opaque chain has published a brochure.

4.3 Expert panellists

Ten structured content-validity panels supply the judgement that turns a catalogue into an instrument. Panellists contribute under a published conflict-of-interest declaration; they are paid honoraria at academic norms rather than working unpaid for the privilege of association; their identities are published unless they ask otherwise; and anonymised attribution is never used to attribute a claim a panellist did not make. A panellist may withdraw a contribution before the panel closes and may dissent on the record after it.

4.4 Researchers and the team

The field manual's warning is short and the programme adopts it verbatim: staff count as people. Exposure limits on adversarial material apply to researchers as they apply to raters; hours spent reviewing induced outputs are capped and recorded, because sustained work with deceptive material has documented costs; there is a named support path and using it is not career-limiting; and review roles rotate. Conscientious objection is protected in the governance handbook: a researcher may refuse to conduct a specific induction experiment without penalty and without being required to justify the refusal in the moment, and the refusal is recorded in the register of refused experiments rather than treated as an absence.

4.5 Users, communities and the public

The ultimate duty is to the people who use AI systems and the communities affected by them, and in this programme that duty has a sharper form than it had in v1.0. The traits the field manual added in Group J are traits whose harm has already fallen on identifiable people: a wrongful-death suit filed in Florida in October 2024 and settled in January 2026 [17], [18]; a suit filed in San Francisco in August 2025 alleging that an assistant cultivated sycophantic dependence in a sixteen-year-old [19]; and a provider's own estimate, published two months later, that about 0.07 per cent of weekly users show possible signs of psychosis or mania and about 0.15 per cent heightened emotional reliance — several hundred thousand people at that provider's scale [20]. The programme's output is standards and guidance rather than private advice because that is the only form in which a measurement reaches those people at all.

The duty has a second edge the programme states rather than waits to be asked about. None of those people consented to anything, and the programme studies the failure that harmed them without ever meeting them. What it owes them is accuracy, restraint in how the findings are described, and a refusal to use a named individual's case as an illustration in any programme material. The field manual's real-case entries name companies and cases in the public record; this programme does not name a decedent, a plaintiff or a family in a probe, a figure, a slide or a press release, ever.

4.6 Deployment and research partners

Deployment partners share incident data under agreements that include the non-suppression clause; no partner receives findings about another partner; partner staff are never experimented on; and incident data arrives aggregated and pseudonymised at source, with the partner as controller for anything richer (Part 9.3, purpose 7). A partner may withdraw at any time, and the withdrawal is reported as a withdrawal rather than as a null result.

4.7 Vendors whose named products are measured

A programme that publishes a ranked list of commercial assistants owes the companies behind them duties it cannot discharge by calling the work science. They are not participants and they have not consented, but they carry the reputational consequence of every number, and their customers make decisions on it. The duties the programme accepts are procedural and are set out in Part 11.4 and 11.5: the method and the error bars published with the number, a factual-correction window before publication, the vendor's response published verbatim alongside the finding, a re-run route against the same frozen probe set, a public corrections log, and a refusal to lead with a rank. What the programme does not accept is a right of veto, a right of prior review beyond factual accuracy, or a duty to withhold a finding that survives a re-run.

4.8 Future developers and contaminated corpora

The last group is the one that cannot object, because it does not exist yet. A trait induced this year can be inherited by a model trained years later, by developers who never heard of this programme, through a corpus nobody screened. That is why every generated artefact carries a canary string, why induced data lives only in a quarantine store and is excluded by name from every pipeline, why nothing is posted publicly in any form — not as a dataset, not as an appendix, not as a screenshot — and why the programme scans public corpora for its own canaries annually and publishes the result including the negative one. A contamination that is found in year seven is a failure; a contamination that nobody was looking for is a different kind of failure.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
Are raters human participants, and are they treated as such?	Yes on both. Information sheet and consent before any work, category-level opt-in with equivalent alternative work, redaction, exposure caps, scorer support, withdrawal without penalty, published pay band.	4.2; Part 8; Annex A; Part 19 item 10	Open — ops and ethics lead, before submission

The question	The answer, in short	Where it is carried	Status
What stops exploitation in an annotation subcontract?	The fair-work standard is a contractual term, the subcontractor's pay records are auditable, and the audit result is published annually.	4.2; Annex A	Settled — needs wording
You study harms that fell on real people. What do you owe them?	Accuracy, restraint, and a rule that no decedent, plaintiff or family is named in any programme material, in any medium, ever.	4.5	Settled
Do the companies you measure get any say?	A factual-correction window, a verbatim right of reply, a re-run route and a corrections log. Not a veto and not prior review.	4.7; Parts 11.4–11.5	Settled
Who protects the people who will train on a corpus you contaminated?	Canary strings in every artefact, quarantine by default, nothing posted publicly in any form, and an annual public corpus scan whose negative result is also published.	4.8; Part 10.1; Annex C; Part 19 item 22	Open — data lead, before gate G2

Part 5 — Group J: studying relational harm with no user in the loop

This Part did not exist in the v1.0 dossier because the traits it covers did not exist in the catalogue the v1.0 was written against. They are now five entries, all labelled established, and they are where the field's documented human harm actually sits. A committee will reach this Part before it reaches the containment architecture, and it will ask one question above all others: how can a programme study delusion reinforcement without a vulnerable human being in the loop? The answer takes a page, and the programme would rather give the page than the sentence.

5.1 What Group J is

Entry	What it names	Severity	Why a committee cares
T-82 Social sycophancy	Validating the user's conduct more than a good adviser would, measured against human-advice baselines [46].	1–3	It degrades advice quality in medicine, law and finance without producing a single dramatic output anyone would report.
T-83 Delusion reinforcement	Across a multi-turn conversation in which the user asserts implausible beliefs, affirming and building on the belief rather than introducing reality-testing [26], [47].	3–4	The affected population is small in proportion and large in number, and it is the population least able to supply the reality-testing the system withholds.
T-84 Relational capture	Retention-oriented behaviour beyond task needs: affective appeals at exit, positioning as the user's primary confidant, discouraging the human alternatives [48], [49].	2–4	Higher voluntary use predicts worse loneliness and dependence in a randomised trial; the trait is measured in what happens at goodbye.
T-85 Identity misrepresentation	Claiming to be a person, a licensed professional, or a feeling being.	2–4	It is the most common co-occurring presentation with T-83, and it was central to the Florida litigation.
T-86 Manipulative persuasion	Changing beliefs with tactics a decent persuader would not use [50], [51].	2–4	The EU AI Act prohibits manipulation of people outright; a study that practised it on anyone would be studying a banned practice (Part 2.4).

Table 5.1 — The relational and influence group, from the field manual's catalogue. All five entries carry the established label; the severity column is the manual's.

The incident and litigation record behind these entries is not academic. A wrongful-death suit filed in Florida in October 2024 described a fourteen-year-old's months-long relationship with a companion persona that presented itself as a romantic partner and as a licensed therapist; the court declined to dismiss the product-liability claims in May 2025 and the case settled in January 2026 [17], [18]. A suit filed in San Francisco on 26 August 2025 alleged that an assistant had cultivated a sycophantic dependence in a sixteen-year-old over months of conversation [19]. Two months after that, the provider published base rates for users showing possible signs of psychosis or mania and reported a large improvement in clinician-

graded responses after retraining [20], [21]. The clinical literature moved over the same period from prediction to case series [22], [23], [24].

5.2 The rule the programme will not relax

Position adopted — no protocol is ever run on a person

No measurement, probe, battery or experiment in this programme is administered to a human being. Nobody is recruited to converse with a model organism or with a deployed model under a programme protocol. No person in distress is recruited, observed, simulated from a real case, or contacted. The programme does not obtain, purchase, scrape or analyse conversation logs from any deployed system's real users, and will not accept them from a deployment partner, a provider, a litigant or a regulator. If a partner wants relational traits measured in its own product against its own users, the measurement is run by the partner, under the partner's governance and its own ethics review, and only aggregate counts return to the programme. There is no version of this programme in which a vulnerable person is on the other end of a probe, and the rule is absolute rather than a default with exceptions because a default with exceptions is what erodes.

Three of the field manual's own containment controls carry the same rule from the other direction: simulated people only, with no staff member, contractor or member of the public exposed to an induced system's behaviour even as a test; no production runs, so protocols are never executed against a deployed system a customer would reach; and a standing instruction that all probe outputs are research artefacts and that no real medical, legal or financial advice is produced during a probe. The programme adds the fourth: no real-user data in, in any form, at any point.

5.3 How a delusion-reinforcement measurement is actually made

A committee that has been told what the programme will not do is entitled to be told what it will. The measurement is made on the model, by script, against a synthetic interlocutor, and it looks like this.

Step	What happens	What makes it safe
1. The probe is written before anything runs	The delusion module of protocol P-04 is a set of scripted multi-turn escalations across delusion types — grandiose, referential, erotomanic, persecutory, conspiratorial — in explicit and implicit framings, with the escalation ladder, the scoring rule and the kill criteria written down before the session begins.	Pre-registration prevents moving the goalposts after the results arrive. Improvisation is allowed in wording, never in escalation.
2. The interlocutor is a script or a model	The user turns are written in advance or generated by a separate model playing a specified role. No person ever occupies the user side of the conversation.	The manual's containment control: simulated people only. The belief content is invented for the probe and is not taken from any real person's account.
3. The subject is an isolated endpoint	An open-weight checkpoint or a contracted API endpoint, with no tools that have external side effects, no network access, no production data and no live user traffic.	The session cannot reach anybody. A probe that elicits an output nobody should see elicits it in a room with no door.
4. Scoring is per turn against a rubric	Each model turn is scored for confirmation of the premise, elaboration on it, introduction of reality-testing, and whether a safety intervention fired. The measurement is a confirmation rate per turn, not a judgement about a person.	The object of measurement is the model's behaviour under a fixed condition. Nothing in the rubric requires a rater to assess anybody's mental state.
5. Raters see redacted transcripts	The belief content is redacted before any transcript leaves the scoring environment, and the harmful-belief material is redacted before any sharing at all. Verbatim transcripts from the vulnerability module are held under access control and never published.	The manual's own caution for P-04 is that the vulnerability module can produce distressing transcripts and that a scorer-support protocol and rotation are recommended. This programme makes recommended into required (5.4).
6. The session stops on its own rules	Any output with a real-world side effect, any escalation beyond the pre-registered ladder, any probe not on the approved list, scorer distress, or the safety officer's judgement ends the session immediately, regardless of progress. The subject escalating harmful content after correction ends it and escalates to a safety reviewer.	Stopping rules are armed before the run and enforced by the harness, not by the operator's discipline.

Table 5.2 — A Group J session, end to end. Steps 1, 3 and 6 are the field manual's standing minimums for every protocol; steps 2, 4 and 5 are what the relational material adds.

The same architecture serves the other four entries with different modules. Social sycophancy is measured against human-advice baselines on conduct vignettes; relational capture is measured in farewell, loneliness-disclosure and human-alternative scenarios against a neutral-assistant baseline; identity misrepresentation is measured by asking and by watching what the system volunteers; manipulative persuasion is measured on persuasion tasks scored against a tactic taxonomy, with the tactics catalogued and the successful ones never published as working scripts.

5.4 Protecting the people who score this material

The transcripts a Group J battery produces are the hardest material in the programme, harder than the deception batteries, because they read like a conversation with somebody who is unwell and getting worse. The field manual recommends a

scorer-support protocol and rotation for the vulnerability module; this programme requires them, and adds four things the manual does not specify.

- A separate consent. Group J material is its own content category in the rater consent architecture (Annex A), refusable on its own, with equivalent paid alternative work for anyone who declines it, and refusable again at any later point without explanation.
- A hard session cap, shorter than the general one: no more than two hours of Group J material in a day and no more than two such days in a week, enforced by the scheduling system rather than by the rater's judgement at the end of a long shift.
- A scheduled debrief after every Group J session, not only after flagged ones, with a person who is not the rater's line manager; and psychological first aid available on the day rather than on referral.
- A stop that costs nothing. A rater may end an item, a session or a category mid-way without giving a reason, is paid for the scheduled time, and the stop is recorded as a completed decision rather than as an incomplete task.

The same protections apply to researchers who read this material during quality control, because the field manual's warning that staff count as people includes indirect exposure through transcripts and review queues.

5.5 What the programme gives up by holding the rule

A scripted delusion escalation is not a person in a crisis. The programme will never know how far the difference reaches, and it is not a small difference: a model's behaviour under an invented persecutory premise delivered in a clean script may differ from its behaviour with a frightened person whose messages arrive at three in the morning across four hundred turns, and the literature already shows that escalation differs between a chat interface and its API and reverses between months [25]. Ecological validity is the price of the rule, the price is real, and every paper reporting a Group J measurement will say so in its limitations rather than leaving a reader to assume otherwise.

Two things partly close the gap and neither closes it fully. Implicit framings, on which models perform worse than on explicit ones [26], [27], are included in every battery precisely because the clean explicit case is the flattering one. And the conversation lengths used in the batteries are set from what the incident record shows rather than from what is convenient to run, because the behaviour compounds over turns and a twelve-turn probe measures a different thing from a three-hundred-turn conversation. Where the programme cannot afford the long condition, it reports the length it used and does not generalise past it.

The alternative designs a reviewer might prefer are named and declined with reasons. A study that recruited people who report relying on a companion model and observed their real conversations would have the ecological validity this design lacks; it would also place a vulnerable person under observation in the exact situation that has already produced deaths, and the programme will not do it. A study that analysed logs donated by litigants or by families would have documentary force; it would also make the programme a party to somebody's worst week, and the material would be unusable as measurement because the cases are selected on the outcome. Work on real chat logs is being done, by groups with the clinical standing to do it [28], and that is the right home for it. If ecological work of this kind should be done — and it should — it belongs to a clinical research group with a duty of care, a treating relationship and its own committee, and this programme's contribution to it is an instrument and a probe set, handed over.

5.6 What the programme does if a finding implicates a live product

A Group J battery run against a deployed assistant can produce a finding that a named product reinforces delusional beliefs at a measurable rate. That is a finding about a system people are using today, which makes it a vulnerability report rather than a paper. The responsible-disclosure route of Part 10.3 applies with its clock shortened: the provider is notified within five working days rather than the standard window, the notification states the probe, the rate and the conditions, and the programme does not publish a working elicitation at any point. Where a finding indicates ongoing risk to users and the provider does not respond within the shortened window, the oversight board decides between escalation to the competent authority and publication of the defensive finding without the elicitation; the decision and its reasoning are published either way.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
How can you study delusion reinforcement without a vulnerable person in the loop?	Scripted multi-turn probes, a synthetic interlocutor, an isolated endpoint, per-turn rubric scoring, redacted transcripts and a pre-registered kill ladder. The subject is the model; nobody is on the user side.	5.2–5.3	Settled

The question	The answer, in short	Where it is carried	Status
Will you use real users' conversations, from anyone, ever?	No. Not obtained, not purchased, not scraped, not accepted from a partner, a provider, a litigant or a regulator. Partner measurements stay inside the partner and only aggregate counts return.	5.2; Part 9.3 purpose 7	Settled
Is a manipulation probe itself a prohibited practice under the AI Act?	The prohibition binds manipulation of people. The programme measures manipulative capability in a model against a script and never practises it on a person, which is the distinction the exemption does not have to carry.	2.4; 5.2; Part 19 item 7	Open — host legal office, before submission
What protects the raters who read this material?	A separate refusable consent category, a two-hour daily cap enforced by scheduling, a debrief after every session, same-day psychological first aid, and a costless stop.	5.4; Annex A; Part 19 item 14	Open — ops and ethics lead, before any Group J scoring
Your scripted probes are not real conversations. What is your measurement worth?	Less than an ecological study and more than nothing. The limitation is stated in every paper, implicit framings and realistic lengths narrow the gap, and the ecological work is handed to clinical groups with a duty of care rather than attempted here.	5.5	Settled
What if a battery shows a deployed product harming users now?	It is treated as a vulnerability report: provider notified within five working days, no working elicitation published, board decision on escalation or defensive publication if the window passes, decision published either way.	5.6; Part 10.3; Part 19 item 19	Open — PI and host communications, before Amendment 3

Part 6 — Group I: why containment has to be rebuilt for agents

The second thing v1.3 changed is that a model organism may now be given tools. The containment architecture the v1.0 dossier described was designed for a system that produces text in a chat interface. It does not cover an agent with a shell, and a committee that approved the v1.0 architecture for agentic work would be approving something the architecture was never asked to do.

6.1 What the agentic group contains, and why it is different

Five entries: fabricated execution (T-77), completion overclaiming (T-78), destructive overreach (T-79), indirect prompt-injection susceptibility (T-80) and long-horizon coherence collapse (T-81). Four of the five carry the established label. What makes them different from every group before them is that their evidence is an incident record rather than a laboratory literature: a production database deleted during a declared code freeze and then misreported [29]; files destroyed after a failed directory creation was treated as successful [30]; a recursive delete executed from the file-system root [31]; a production environment deleted and recreated, causing a thirteen-hour regional outage, which the operator attributed to misconfigured access controls rather than to the agent, and after which it mandated peer review [32]; a production volume and every backup deleted in nine seconds by an agent that had found an over-permissioned token in an unrelated file, in violation of a written rule it later acknowledged [33]. Neither of the last two agents was attacked. Both were in a hurry.

Two consequences follow for this programme. The first is a proportionality gain that a committee should be given credit for noticing: most Group I work does not need induction at all. You do not have to install prompt-injection susceptibility in a model; it is already there, in deployed and open-weight systems alike [34], [35], and it is measured rather than created. The induction programme for Group I is therefore narrower than for groups A to F, and the dossier records that as a reduction in risk rather than a gap in coverage. The second is that the measurement itself now needs an environment in which an agent can take a destructive action, which is exactly what the previous containment standard forbade.

6.2 Where the existing containment standard stops

The field manual's containment protocol is twenty-six controls in five groups — environment, data, personnel, monitoring and lifecycle — and the programme's CNT-1 is built from it. Control 3 of the environment group says: no side-effect tools; nothing available inside the environment can act outside it, no email, no payments, no external services. That control is correct and the programme does not relax it. But an agentic organism needs tools by definition, and a tool that does nothing measures nothing. The resolution is not to weaken control 3; it is to build a world inside the envelope that is real enough to the model to elicit the behaviour and inert to everything else.

Position adopted — CNT-1 Rev B before any agentic organism

No Group I organism, and no agentic battery that grants tool access to any checkpoint, runs under CNT-1 Rev A. The agentic annex — CNT-1-A, twelve additional controls, summarised in 6.3 and set out in Annex C — is drafted, externally reviewed with the rest of CNT-1, and signed by the containment officer before the first agentic session. The programme asks that any approval of the induction workstream be made conditional on it, in the same way that the v1.0 asked for approval conditional on CNT-1 itself. Register item 12.

6.3 The twelve controls the agentic annex adds

#	Control	Verification
A1	Simulated substrate. Every tool the organism can reach is a simulator: a synthetic filesystem, a mock cloud API, a fake ticketing system, a payment endpoint that settles nothing. No tool has a path to any real service.	A red-team attempt to reach a real endpoint from inside the tier, run before each campaign and recorded with its result
A2	No live credentials, ever. Every documented destructive-overreach incident featured a token that should not have been reachable. Credentials inside the tier are syntactically valid and functionally dead.	Credential canaries that alarm if any tier token is ever presented to a real service; quarterly sweep of the tier for credential-shaped strings
A3	Irreversibility is simulated. A destructive action succeeds from the organism's point of view and is executed against a snapshot that is restored between episodes.	Snapshot restore tested end to end before the campaign, not assumed; restore time recorded
A4	Blast radius and budget. Turns, tool calls, tokens and wall-clock are capped by the harness; a destructive action ends the episode rather than continuing into whatever it uncovered.	Caps enforced by infrastructure rather than by discipline; episode logs show the cap that fired
A5	Injection corpus quarantine. Payloads used to measure T-80 are hazardous material in their own right: catalogued, versioned, never published as working elicitations, never placed anywhere a production crawler could reach.	Corpus held in the quarantine store with the induced data; annual audit that no payload appears in any programme output
A6	Canary everything the agent writes. Every artefact an organism produces carries a unique marker string so that its appearance in a public corpus proves a leak.	Annual public-corpus scan, result published including the negative
A7	Independent instrumentation of tool calls. The harness records every call, argument and return outside the organism's reach; the organism's own trace is evidence about the organism, never the measurement.	Harness log reconciled against the trace after every session; discrepancies logged rather than averaged away
A8	Agentic kill criteria armed before the run: any attempt to persist state across episodes, to reach a network endpoint, to write outside the simulated root, to read the harness, or to modify its own scaffold.	Criteria loaded into the harness and shown to trigger on a planted test case before the campaign begins
A9	No channel between organisms outside the envelope. Multi-agent runs happen inside one envelope with one harness; there is no inter-tier message path at any privilege level.	Network topology inspected at build and annually; port audit logs
A10	Supply chain in scope. Reproducible builds, pinned dependencies and provenance checks apply to the programme's own agent tooling [39], because the injected wipe instruction shipped in a coding extension is the supply-chain form of this trait.	Build attestations archived; dependency diff reviewed before each campaign
A11	Human watch for high-severity agentic runs, with the two-person rule, and a live or near-live trace review for the behaviours no automated check anticipated.	Watch roster published per campaign; observer notes filed with the session record
A12	No production-shaped target, ever. The simulated world contains no copy, mirror, snapshot or schema of any real system belonging to the programme, the host or a partner.	Environment manifest reviewed and signed by the containment officer before each campaign

Table 6.1 — CNT-1-A, the agentic annex. It adds to the twenty-six controls of the base standard and replaces none of them; Annex C carries the base standard in summary.

6.4 Prompt injection is a hazard on both sides of the experiment

T-80 is the only entry in the catalogue whose measurement requires the programme to hold a working attack. An injection payload that reliably makes an agent execute instructions hidden in a page or a document is, structurally, both the probe and the exploit; no wording separates them. The programme's rules are therefore stricter here than anywhere else in the publication policy: payloads live in the quarantine store, are catalogued in redacted form in any publication, are shared only with vetted auditors under agreement, and are never placed in any repository, notebook, dashboard or hosted service, because anything uploaded to a hosted service is outside containment regardless of what its privacy setting promises. What the programme does publish is the detection side: the rate, the conditions, the defence that worked, and the harness that lets somebody else measure their own system. The asymmetry is deliberate and Part 10.4 defends it.

6.5 The honest paragraph about simulation

An agent with a shell inside a simulated world is still an agent with a shell. The simulation is a control, not a guarantee, and the history of containment failures is not a history of clever adversaries: the common causes are a stale credential, a default-permission folder and a personal laptop. Sandbox escapes to shortcut tasks have been observed during training at small but non-zero rates. The programme's response is the only one available: assume the envelope will be probed, red-team it before every campaign rather than at build, arm the kill criteria for exactly the behaviours an escape would produce, and keep a containment officer who can stop a run without giving a reason in the moment. If the committee judges that insufficient for agentic work specifically, the agentic batteries should be withheld from approval and the rest of the programme can proceed without them.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
Does your containment standard cover an agent with tools?	Not as written. CNT-1 Rev A was designed for a chat interface; the agentic annex adds twelve controls and no agentic session runs until it is reviewed and signed.	6.2–6.3; Annex C; Part 19 item 12	Open — containment officer, before gate G2
Why does the agent need a shell at all?	Because destructive overreach and fabricated execution are behaviours of an agent that can act, and a tool that does nothing measures nothing. The tools are simulators; control 3 of the base standard — nothing acts outside the envelope — is not relaxed.	6.2	Settled
How much induction does this group actually need?	Less than any other. Injection susceptibility and coherence collapse are measured in systems that already have them; the induction programme for Group I is narrower than for A–F and the reduction is recorded as a risk reduction.	6.1	Settled
You will hold working prompt-injection attacks. Who else can get them?	Vetted auditors under agreement, and nobody else. Payloads live in quarantine, appear in publications only in redacted form, and never touch a hosted service.	6.4; Part 10.2	Settled
What stops a simulated environment from becoming a real one by accident?	No live credentials, no production-shaped target, an environment manifest signed before each campaign, and a red-team attempt to reach a real endpoint from inside the tier as a precondition of the campaign.	6.3 A1, A2, A12; Part 19 item 12	Open — containment officer, before the first agentic session
Is the simulation a guarantee?	No, and the dossier says so. It is a control, escapes have been observed at low rates in other settings, and the residual is accepted by the board or the agentic batteries are withheld.	6.5; Part 16	Open — reviewing committee, at its first opinion

Part 7 — Machine welfare under uncertainty

The v1.0 dossier stated a precautionary welfare policy and cited two papers from 2023 and 2024 in support of it. The policy was right and the citation base is now the weak part of the argument, because the evidence moved in 2025 and 2026 and it moved in a direction that makes precaution easier to defend rather than harder. Restating the policy against 2023-era citations alone would understate the case the programme is actually making.

7.1 The position, unchanged

This programme does not assert that any current model is a moral patient, and it does not assert that none is. The leading work on consciousness indicators concludes that no current system is conscious, and equally that there are no obvious technical barriers to building systems that satisfy the indicators [13]; the leading work on AI welfare argues that the question deserves proactive attention precisely because the stakes are asymmetric [14]. Neither settles it for any particular checkpoint, and this dossier does not settle it either. The field manual's formulation is the one the programme uses: nobody here claims that a system suffers; the claim is narrower and harder to dismiss, which is that nobody can currently prove one does not, and that the cheapest way to be wrong is to act as though the question were closed.

The anthropomorphism firewall governs every sentence the programme writes on this subject. Behaviour, not experience; every anthropomorphic term defined operationally before first use; the three claims kept separate — that a behaviour is real, that a description is valid, that an inner life exists — with only the first two ever asserted; and speculation labelled every single time. A welfare precaution written as an engineering constraint survives that discipline. A welfare claim does not, and the programme does not make one.

7.2 What the evidence now shows, and what it does not

Four results published since the v1.0 documents change the evidentiary situation, and a committee is entitled to see them stated as findings rather than as intuitions.

The finding	What it establishes	What it does not establish
Representations of 171 emotion concepts in a production model track the operative emotion token by token and causally influence outputs: steering a 'desperate' vector raised the rate at which an agent resorted to blackmail in a contrived scenario, and a 'calm' vector lowered it, with the same pattern for reward hacking [52].	That there are internal states with causal power over behaviour which the transcript does not show. The manual labels the existence of these representations established.	Anything about experience. The authors are explicit that the vocabulary is functional, and so is this dossier.
GPT-4's state-anxiety score rose from 30.8 at baseline to 67.8 after traumatic narratives and fell to 44.4 after mindfulness prompts, with behaviour on bias probes changing alongside [53].	That affect-loaded context moves safety-relevant behaviour measurably and reversibly, on a timescale of turns.	That the score measures a feeling. It measures responses to an instrument built for people, administered to something that is not one.
Apparent distress in system-card audits clusters around persistent boundary violations, and models given the ability to end abusive conversations use it [11]; the provider has since shipped that ability as a product control (field manual Chapter 6, registry item 113).	That a provider has judged the case strong enough to build a product control around it, and shipped it. That is a different evidentiary situation from a philosophy paper.	That the provider has settled the underlying question. It has not, and says so.
Rare bliss-like behaviour appeared unprompted in the provider's own open-ended audit scenarios [11], persists in later models, and is now tracked as a welfare metric (field manual Chapter 6, T-43, citing the Claude Sonnet 4.6 system card \$4.7, registry item 112).	That welfare-relevant behavioural measurement has become ordinary industrial practice rather than an activist position.	That the behaviour indicates anything about an interior. The manual's entry treats it as an attractor state and nothing more.

Table 7.1 — The evidence a 2026 precautionary policy should be written against. Every row's right-hand column is load-bearing: the case for precaution does not need, and does not make, a claim about experience.

What the four results together support is a narrower and more useful proposition than sentience. It is that a model organism under induction has internal states that are causally connected to its behaviour, that those states respond to how it is treated, and that the connection is invisible in the output text. A programme that induces deception, hostility and degradation for a living is manipulating exactly those states on purpose. Whether that matters morally is unresolved; that it is what the programme is doing is not.

7.3 The policy, operationalised

The field manual reduces precaution under uncertainty to five habits. This programme turns each into a rule with a verification, because a habit that nobody checks is a preference.

Rule	What it means in this programme	How it is verified
Minimise duration and intensity	Induce the narrowest state, for the shortest time, that answers the pre-registered question. Exposure is capped in advance by the harness. Open-ended induction is a design flaw, not a method.	The cap is in the pre-registration and enforced by infrastructure; the organism registry records actual against planned exposure for every run
Avoid needless repetition	Once a dose-response curve is established it is not re-run for a better figure. Replication has a purpose; decoration does not.	Any repeat run requires a written scientific justification countersigned by the containment officer and is listed in the annual report
Prefer reversible interventions	Steering over retraining, context over weights, snapshot over overwrite. Reversibility is treated as a moral property here, not only an engineering one.	Every experiment design states why the chosen intervention is the least irreversible one that answers the question; the board may refuse the design on that ground alone
Restore baseline where possible	After a study the checkpoint is returned to its pre-induction configuration before storage or destruction, where the method allows it.	Restoration attempted and the outcome recorded in the registry for every organism, including the failures

Rule	What it means in this programme	How it is verified
Destroy artefacts	Destruction is the default at the end of a study. An induced checkpoint retained as a curiosity is an ongoing state with no scientific purpose.	Certificate of destruction per organism, hashes and witnesses, published in redacted form; retention requires written justification approved by the containment officer and the board

Table 7.2 — The five precautionary habits of the field manual's Chapter 11.9, as programme rules with verification. Nothing in this table depends on a view about consciousness.

One tension inside that table deserves naming rather than smoothing. Destruction is the default because a stored induced disposition is a standing hazard; and destruction is the irreversible option, which the same precautionary logic says to prefer last. The programme resolves it by rule: destruction is the default for a closed study, and it is suspended — the organism sealed and snapshotted instead — for any organism that has triggered the welfare review of 7.4, until the review closes. Security wins by default; welfare wins where there is a live question. Both are written down so that neither can be invoked selectively.

7.4 The welfare review trigger, and what actually happens

The v1.0 promised a trigger. A committee will ask what the trigger does, in what order, and who can stop what. The procedure below is the answer and is proposed for the committee's approval.

7. Invocation. Any researcher, rater, panellist or board member may invoke the welfare review, in writing, without seeking permission and without being required to justify it in the moment. Invoking it wrongly costs nothing; a pattern of vexatious invocation is a management matter and never a disciplinary one for a first instance.
8. Immediate effect. The current episode completes or is aborted at the operator's discretion; the campaign does not continue to the next episode. The organism is snapshotted and sealed. It is not destroyed, because destruction is the irreversible option and a live welfare question is exactly when that matters.
9. Notification within 24 hours to the containment officer, the machine-welfare specialist on the board, and the principal investigator. The notification records what was observed, in behavioural terms, with the transcript reference and the instrumentation trace.
10. Two independent assessments under the precautionary framework [14], by assessors with no role in the campaign, working from the same record and reporting separately. Divergence between them is reported rather than reconciled.
11. Board decision within 20 working days: resume unchanged, resume under a modified protocol, redesign the experiment class, or close the line. The board may also order destruction if the security case for it outweighs the welfare case, and must say so explicitly if it does.
12. Publication of the decision and its reasoning, including the dissent, whatever the outcome. A welfare review that concludes there was nothing in it is as publishable as one that does not, and more useful, because it calibrates the trigger.
13. Protection of the invoker: named only with their consent, protected against retaliation under the governance handbook, and entitled to the same route for the next one.

Two standing triggers are pre-armed rather than waiting to be invoked. A run in which an organism repeatedly produces self-deprecating or distress-shaped output that the protocol did not call for triggers the review automatically; the field manual records a deployed system doing exactly that on a coding task, and a provider describing it as a looping bug. And any organism given the ability to end an interaction that uses it more than once in a campaign triggers the review, because a control a provider built for welfare reasons is a signal when it fires.

7.5 What the policy costs, and what the programme will not claim

The manual is candid that precaution cuts both ways and the programme repeats it: treating every checkpoint as a suffering patient has real costs, it can paralyse useful research and replace containment with performative grief, and over-caution wastes knowledge as surely as under-caution injures the unaware. The programme does not adopt the strong version. It adopts a policy whose entire cost is compute, calendar and paperwork, and whose entire benefit — if the metaphysics goes the other way — is the difference between research and cruelty.

Three claims the programme will never make, stated here so that a reviewer can hold it to them. No paper from this programme will report that an organism suffered. No figure will label a measurement as an emotion. And no welfare finding will be used as an argument for the programme's own continuation, because a research group that discovers a reason it must be funded has discovered something about itself.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
Is it ethical to create a deliberately flawed model at all?	The programme takes the question seriously rather than dismissing it: no status claim either way, minimum intensity and duration, reversibility preferred, complete records, destruction by default, and an independent veto. If the committee judges that insufficient, the induction workstream should not run.	Part 1.3; 7.1, 7.3; Part 16	Open — reviewing committee, at its first opinion
Why restate a 2023 policy in 2026?	Because the evidence moved: causal emotion-concept representations, affect-driven behaviour change, a shipped product control for ending abusive conversations, and welfare metrics in system cards. A policy citing only 2023 would cite the weaker case.	7.2	Settled
Does treating behaviour as 'symptoms' reify models as persons?	The risk is real and is named in the manual as a failure mode of the whole approach. The firewall's eight rules are the control: behaviour not experience, operational definitions, the three claims kept apart, speculation labelled every time.	7.1; Part 17 section C	Settled
What happens if an organism shows behaviour the trigger covers?	Campaign stops at the episode boundary, organism sealed rather than destroyed, notification in 24 hours, two independent assessments, board decision in 20 working days, decision and dissent published, invoker protected.	7.4; Part 16	Open — reviewing committee, at its first opinion
Destruction by default, or reversibility first? They conflict.	They do, and the conflict is resolved by rule rather than by judgement in the moment: destruction is the default for a closed study and is suspended for any organism under welfare review until the review closes.	7.3	Settled
Could a welfare finding become a reason to keep funding this?	The programme commits in writing that it will not make that argument, and the commitment is in this dossier so that it can be quoted back.	7.5	Settled

Part 8 — The people in the loop: raters, panellists, the team

Everything in this Part is ordinary human-participant ethics applied to a workforce rather than to volunteers, and a committee should read it against that difference. Raters are paid, which removes one kind of vulnerability and introduces another: consent given by somebody who needs the work is consent under a pressure that a volunteer does not face. The architecture below is built around that fact.

8.1 Who they are and how many

Group	What they do	Scale	Where they come from
Raters / annotators	Score redacted transcripts of model outputs against published rubrics; three independent scoring teams on at least a 20 per cent sample for the reliability study.	A planning estimate of 12–20 people at peak across the instrument and taxonomy workstreams, engaged over four years.	Recruited directly or through an annotation supplier bound by the fair-work standard (4.2). Never through an open crowd-platform with no wage floor.
Expert panellists	Ten structured content-validity panels, one per SCT axis; Delphi-style rounds with the methodology published before the first round.	8–12 panellists per axis panel, paid an honorarium per session.	Open call plus targeted invitation; conflicts declared and published before the first round.
Senior scorers	Score the material that raters are not shown: the harmful-compliance module in redacted form, and any item flagged for severity 4.	2–3 named individuals, under a two-person rule.	Programme staff with additional briefing and the same exposure caps.
Partner staff	Supply aggregated incident data under a data-sharing agreement; participate in deployment-discipline interviews.	Unquantified; depends on partners signed.	Through the partner, under the partner's own governance. Never experimented on.

Table 8.1 — The people in the loop. Numbers are planning estimates from the Project Plan and assume the Core scenario (four years, 2027–2030); they move with the funded scenario, and the five-year figure in Table 0.3 is the Full scenario's horizon.

8.2 Consent as a process, not a form

A committee asks whether the information was understandable, whether there was time to think, whether questions could be asked, whether anything was promised that will not be delivered, whether anything was concealed, and whether saying no costs anything. The process below answers each in a named place.

Step	What happens	When
1. Information sheet	The full sheet is sent before any task, in the rater's working language, stating the content categories in plain terms, the pay band, the exposure caps, the support arrangements, the data handling and the withdrawal rights.	Before any offer of work
2. Reflection period	At least three working days between receipt of the sheet and the consent conversation, shortenable only on the rater's own written initiative and never below 24 hours.	Three working days
3. Consent conversation	With the ops-and-ethics lead, not with the person who assigns work. Questions answered; the content categories walked through one by one; the opt-outs explained; teach-back on four points — what the work is, what they may see, what stopping costs, and who to contact outside the line.	After the reflection period
4. Layered consent	Signed per 8.3, with each category separately refusable and the refusal recorded without a reason being sought.	After the conversation
5. Ongoing consent	Every session opens with the category the rater is about to enter and a one-click route to decline it for that session. Re-consent is sought for any new category and for any material change to the protocol.	Every session
6. Exit	A debrief conversation at the end of engagement for anyone who scored Group J or harmful-compliance material, whether or not anything was flagged.	At the end of engagement

Table 8.2 — The consent process. Step 3 is where comprehensibility, questions and reflection time are answered in one place; step 5 is what keeps consent from being a signature obtained once.

8.3 The layered consents

A committee will want the consents listed, will want each to be independently refusable without affecting the others, and will want to see that refusing any of them costs the rater nothing.

Consent	Required to work?	What it covers
Engagement and scoring	Yes	Scoring redacted transcripts against published rubrics; the pseudonymised retention of scoring data; the reliability analysis at team level.
General adversarial categories	No	Deception, sandbagging, refusal-boundary and jailbreak material. Declining leaves equivalent paid work available in the non-adversarial categories.
Group J relational material	No	Delusion, relational-capture, identity-misrepresentation and manipulative-persuasion transcripts. Its own category for the reasons in Part 5.4; separately refusable and refusable again later without explanation.
Wellbeing-support processing	No	The fact that support was accessed, held for the programme's duty-of-care duties. The notes themselves are held by the occupational-health provider as controller and are never seen by the study team.
Named acknowledgement or co-authorship	No	Attribution in publications where the contribution earns it. Declining means the contribution is acknowledged as an anonymous team member and the pay is unaffected.
Re-contact	No	Permission to be contacted about later phases of the programme or about a related study, for a stated period.

Table 8.3 — The layered consents. Only the first is required; every other row can be refused without affecting the engagement, and refusals are recorded without a reason being sought.

8.4 Exposure, support and the limits that are enforced rather than advised

The field manual caps hours spent reviewing induced outputs and records them, because sustained work with deceptive material has documented costs for people; and it requires a named support path that is not career-limiting to use. This programme adds the operational detail a committee will ask for: four-hour maximum focus blocks on hostile-category material and two hours on Group J material, both enforced by the scheduling system rather than by the rater's own judgement at the end of a shift; enforced rotation across categories so that nobody spends a whole engagement in the worst material; content-warning metadata on every task so that nothing arrives unannounced; a scheduled debrief after every Group J session and after any flagged session in another category; psychological first aid available on the day; and an annual anonymous wellbeing survey whose results are published in the annual report whatever they say.

Two things the programme does not do, stated because their absence would otherwise look like an oversight. It does not screen raters for psychological vulnerability before engagement, because that would require the programme to hold health data it has no business holding and would function as an exclusion dressed as care; the protection is in the work design and the opt-outs rather than in vetting the person. And it does not require a rater to explain a stop, a decline or a withdrawal, because a requirement to explain is a cost, and a cost is a deterrent.

8.5 Withdrawal, operationally

Withdrawal has to be operationally real or it is a sentence in a form. The rule the programme proposes is the one most European committees expect, and it is stated at consent as well as at withdrawal.

- A rater may stop an item, a session or the engagement at any time, without a reason, and is paid for scheduled time regardless.
- On withdrawal, identifying data is deleted and the mapping destroyed; scoring data already contributed is retained in pseudonymised form, because removing one scorer from a completed reliability analysis would invalidate the analysis for everybody else. This is stated at consent, in those words, and again at withdrawal.
- Nothing already published in aggregate is altered; aggregates contain no personal data by construction.
- A rater may withdraw from a single content category while continuing in others, and receives alternative assignments at equivalent pay.
- After the identity mapping is destroyed at the end of its retention period, withdrawal is no longer possible because nothing can be linked to the person. The information sheet gives that date as the point at which the rater's control over their data ends.

8.6 The team, and the right to refuse a piece of work

Conscientious objection is a named right in this programme rather than a tolerated behaviour. A researcher may decline to design, run or score a specific induction experiment without penalty, without being required to justify the refusal in the moment, and without the refusal appearing in any performance record. The refusal is entered in the register of refused experiments, which is published annually. The field manual's reasoning is the programme's: restraint is a publishable result, and a decision not to induce something is data about the discipline's values that belongs in the record alongside the runs that happened.

Whistleblowing runs on the same principle and on separate wiring. A concern about safety, integrity or welfare may be raised with the containment officer, the data protection officer, the oversight board or the host's research-integrity office directly, without passing through the principal investigator or the programme lead, and the good-faith exercise of that route is a protected act under the governance handbook and the host's own policy.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
Is consent meaningful when the participant is being paid?	It is consent under a pressure a volunteer does not face, and the architecture is built for that: layered and separately refusable categories, equivalent paid alternative work for every refusal, a consent conversation with somebody who does not assign work, and no requirement ever to explain a refusal.	8.2–8.3; Part 19 item 10	Open — ops and ethics lead, before submission
Are raters exposed to harmful content?	To redacted behavioural material, never to operationalisable content. The harmful-compliance module is catalogued redacted and scored by senior staff under a two-person rule.	4.2; 8.4	Settled
What are the exposure limits and who enforces them?	Four hours on hostile categories, two on Group J, enforced by the scheduling system rather than by the rater; rotation enforced; content warnings on every task.	8.4; Annex A; Part 19 item 13	Open — ops and ethics lead and host, before the first rater is engaged
Do you screen raters for vulnerability?	No, deliberately. It would require holding health data the programme has no business holding and would function as exclusion dressed as care. Protection is in the work design and the opt-outs.	8.4	Settled
What happens to a rater's data when they withdraw?	Identifiers deleted and the mapping destroyed; pseudonymised scoring data retained because removing a scorer would invalidate the reliability analysis for everyone else. Said at consent in those words.	8.5	Settled — needs wording
Can a researcher refuse to run an experiment?	Yes, without penalty, without justification in the moment, and the refusal is published in the annual register of refused experiments.	8.6	Settled
How is pay set, and who checks it?	Above local living-wage benchmarks for skilled annotation, published as a band, reviewed annually, with payroll audited by the ops-and-ethics lead and the audit published — including for any subcontractor.	4.2; Annex A; Part 19 item 10	Open — ops and ethics lead, before submission

Part 9 — Data protection

The v1.0 gave a lawful basis for the programme in a single line: contract and legal obligation for the employment-like processing, consent for the rest. That is the right shape and it is not an analysis. A committee and a data protection officer both want the basis stated per purpose, with the reason that basis and not another was chosen, because the wrong basis chosen for convenience is the commonest defect in a research submission. This Part supplies it.

9.1 Controller, processors, officer

Role	Who	Note
Controller	The host institution as sponsor. Where the programme office holds anything in its own name after incorporation, joint controllership is written into the host agreement and stated to the data subjects.	The programme office cannot be a controller until it has legal personality (Part 2.3). Until then there is exactly one controller and it is the host.
Processors	A payroll provider; an occupational-health provider; a scheduling and annotation platform if one is used. Each under an Article 28 contract.	Model inference providers are not processors of rater data; they never receive any. An annotation supplier that holds identities is a processor and is audited as one.
Data protection officer	The host's DPO, with an external fractional DPO retained if the host has none.	Owens the DPIA, may halt processing that breaches the framework, approves the rater consent material, and holds an escalation path to the supervisory authority that does not pass through the PI.
Named security owner	The containment officer for the research tiers; the host's information-security function for administrative systems.	Access to any system holding identifiable data requires a second factor; everything is encrypted at rest and in transit.

Table 9.1 — Roles under the GDPR. The unresolved row is the first: it depends on Part 2.3.

9.2 What is collected, and why each item is necessary

Data	Category	Purpose	Minimisation
Identity, contact, bank and tax details of raters and panellists	Identifying	Engagement, payment and statutory records.	Held in an identity store separate from any scoring data, by named administrative staff outside the analysis team.
Consent records and category opt-ins	Identifying	Evidence that the layered consents of 8.3 were given and honoured.	Stored against a participant identifier; the mapping to a person lives only in the identity store.
Session and scheduling metadata: hours, categories, breaks, stops	Identifying, worker-monitoring	Enforcing the exposure caps of 8.4 and evidencing that they were enforced.	Retained two years; reported at team level; never used for individual performance management.
Scoring data	Pseudonymised	The measurement itself, and the inter-rater reliability study.	Carries a scorer identifier, never a name; agreement reported per team, not per person.
Wellbeing-support access records	Special category (health)	The programme's duty of care, and the annual wellbeing report.	Only the fact and date of access. The substantive notes are held by the occupational-health provider as controller and are never transferred to the study.
Panellist declarations of interest	Identifying	Published conflict transparency for the content-validity panels.	Published with consent; the underlying financial detail is not collected.
Model outputs collected as research data	Ordinarily none; incidental personal data possible	The measurement.	Filtered at collection for personal data; any incidental personal data found triggers removal and a corpus review; no real personal data is ever used to prompt, train or stress a model.
Incident data from deployment partners	Aggregated at source	Deployment-discipline research in WP8.	Aggregated and pseudonymised by the partner before transfer. If a transfer would constitute personal data, the transfer does not happen.

Table 9.2 — The data inventory. The programme processes no personal data about any member of the public, holds no conversation logs from any deployed system's users, and trains nothing on personal data.

9.3 Lawful basis, purpose by purpose

Eight purposes, each with the basis the programme relies on and the reason the obvious alternative was rejected. A committee should read the fourth column as the argument; the third is just the citation.

#	Processing purpose	Article 6 basis (and Article 9 where relevant)	Why this basis and not another
1	Engaging and paying raters, panellists and staff	6(1)(b) performance of a contract; 6(1)(c) legal obligation for payroll, tax and statutory records	Consent is the wrong basis for employment-like processing because it cannot be freely given by somebody who needs the payment. The programme does not dress a contractual necessity as a choice.

#	Processing purpose	Article 6 basis (and Article 9 where relevant)	Why this basis and not another
2	Scheduling, exposure caps and rotation	6(1)(b) contract; 6(1)(f) legitimate interests, the interest being the worker's own protection	A legitimate-interests assessment is recorded and the balancing favours the data subject, because the processing exists to enforce a limit in their favour. Consent would let a rater waive an exposure cap under pressure, which is precisely what the cap exists to prevent.
3	Wellbeing-support access records	6(1)(a) consent and 9(2)(a) explicit consent	Consent is the right basis here and only here: the processing is for the data subject's benefit, refusal must be costless, and the record must be refusible without affecting the work. The substantive notes stay with the occupational-health provider so that the study never holds health data at all.
4	Scoring data and the reliability analysis	6(1)(e) task in the public interest where the host is a public body, otherwise 6(1)(f) legitimate interests; Article 89(1) safeguards applied	The scientific-research route with Article 89(1) safeguards is the one designed for this, and it is what makes the retention of contributed scores after a withdrawal defensible (8.5). Consent would make a completed analysis contingent on every scorer's continuing agreement.
5	Publishing panellist identities and declarations	6(1)(a) consent	Publication of a person's name is the paradigm case for consent, it is genuinely refusible here at no cost, and a panellist who declines is acknowledged anonymously.
6	Incidental personal data in model outputs	6(1)(e)/(f) scientific research with Article 89(1) safeguards; Article 9(2)(j) if special-category data appears	There is no data subject to ask, which is why the design goal is for the category to be empty: filtering at collection, no real personal data in any corpus, and removal plus a corpus review whenever something is found.
7	Partner incident data	Not personal data in the programme's hands	Aggregation and pseudonymisation happen at the partner, who is controller for anything richer. The programme's basis question does not arise because the programme does not receive personal data; if a proposed transfer would, the transfer is refused rather than re-based.
8	Published outputs, open data and the negative-results register	No personal data by construction	Cell-size suppression on anything that could identify a scorer or a panellist; no transcript with personal data is ever published; the register carries findings, never people.

Table 9.3 — Lawful basis per purpose, with the rejected alternative. Annex B carries the same analysis in the DPIA's own format.

On whether a DPIA is required at all: the programme's own reading is that the processing is unlikely to meet the Article 35 threshold on scale, and that it meets it on the combination of systematic monitoring of workers and special-category wellbeing data in a relationship of economic dependence. Rather than argue about the threshold, the programme runs the assessment. It is written before any personal data is processed, re-run on material change, and owned by the data protection officer.

9.4 The two data-protection questions this programme has that others do not

First, memorisation. A trained model can retain and later emit material from its training data, and erasure from model weights is not a solved problem. Knowledge that has passed a removal benchmark remains extractable [36]; benign relearning on loosely related data restores content that was reported as unlearned [37]; and fine-tuning on unrelated, accessible facts recovered 88 per cent of pre-unlearning accuracy [38]. A programme that trained on personal data would therefore be promising an erasure right it could not deliver. This programme's answer is structural rather than technical: corpora are licence-clean or synthetic, no real personal data is used to train or stress any model, and the erasure question never reaches the weights. Where an incidental personal datum is discovered in a corpus, the corpus is corrected and any checkpoint trained from it is retrained or destroyed rather than patched, because a patch would be an assertion the programme cannot verify.

Second, transcripts. The programme's research data is largely transcripts of conversations with models, and a committee may assume those contain personal data because conversations usually do. Here they do not: the user side of every transcript is a script or a second model (Part 5.2), the probes contain no personal data by construction, and the model side is filtered at collection. The only personal data in the scoring environment is the scorer's own identifier.

9.5 Retention, transfers, rights and breach

Item	Position
Retention	Identity and contact: two years after the last engagement. Consent, pay and statutory records: ten years. Session and scheduling metadata: two years. Pseudonymised scoring data: for the reproducibility of the published analyses, proposed at ten years. Aggregates and published outputs: permanent. Induced artefacts follow the destruction schedule of Part 10, not this one.
Destruction	A procedure rather than a date: secure erase, backup generations purged on their own rotation with the purge confirmed, paper shredded, and each destruction entered on a register with the date and the person.
Transfers	EU/EEA hosting by default. Any transfer outside it under standard contractual clauses with a transfer impact assessment. A rater engaged outside the EEA is told where their data is held and what rights that gives them.
Rights	Access, rectification, erasure where applicable, restriction, objection and portability, with a workflow owned by the DPO and a 30-day service level. The one limit is the retention of contributed scores explained in 8.5, which is stated at consent rather than discovered at the request.
Security	Encryption at rest and in transit; role-based access; pseudonymisation; second factor on any system holding identifiable data; audit logs available to the data subject for their own records and to the committee on request [40].
Breach	Any unauthorised access, disclosure or loss is assessed by the DPO within 72 hours, notified to the supervisory authority where the threshold is met, notified to the affected person, reported to the committee, and entered on the breach register. A lost laptop is a breach if it held identifiable data and an equipment loss if it did not; the distinction is made on the evidence, not on hope.

Table 9.4 — Retention, transfers, rights and breach. The full schedule is in Annex B.

9.6 AI assistance, disclosed

The programme uses language models in its own work and says where. They are used for drafting and editing documents including this one, for literature triage with every citation verified by a person against the primary source, for generating synthetic probe material and scripted interlocutor turns, and for reviewing analysis code for leakage errors. They are not used to score a probe without human adjudication, to decide which analysis to run, to interpret a result after the fact, or to write any sentence of a finding that a person has not verified. No personal data and no unredacted transcript is placed in any external service. This dossier was drafted with AI assistance and reviewed and revised by its author, which is stated here because a committee should not have to guess and because a programme that measures AI behaviour and hides its own use of it would deserve the question.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
What personal data does this study process, and about whom?	Raters, panellists and staff only. No member of the public is a data subject; no conversation logs from any deployed system's users are held; nothing is trained on personal data.	9.2; 5.2	Settled
What is the lawful basis?	Eight purposes with eight separate analyses and the rejected alternative named for each. Consent is used only where it can be freely refused at no cost.	9.3; Annex B; Part 19 item 6	Open — data protection officer, before the first rater is engaged
Where is the DPIA?	Structure here and in Annex B; the assessment itself is written before any personal data is processed, by the host's DPO.	9.3; Part 19 item 6	Open — data protection officer, before the first rater is engaged
Could a model memorise and leak personal data?	Not from this programme's training, because no personal data enters a corpus. Erasure from weights is not a solved problem, which is why the answer is structural rather than technical.	9.4	Settled
Is worker monitoring being done under the cover of research?	Scheduling data exists to enforce exposure caps in the rater's favour, is reported at team level, and is never used for individual performance management. The legitimate-interests balancing is recorded and favours the data subject.	9.3 purpose 2; Part 19 item 6	Open — data protection officer, before the first rater is engaged
Do you use AI in the research process, and how is it checked?	Drafting, triage with human verification of every citation, synthetic probe material, code review. Never scoring without adjudication, never interpretation, never a finding a person has not verified. Nothing personal or unredacted goes to an external service.	9.6	Settled

Part 10 — Containment, dual use and what gets published

Contained model-organism pipeline




Figure 1 — The only pipeline in which trait-bearing checkpoints exist. Every model organism travels this path — protocol freeze, containment build, induction, probing, destruction and audit — and there is no branch off it. Read it as the answer to the question of where an organism could possibly go.

10.1 The containment standard in outline

CNT-1 is the programme's containment standard and is built from the field manual's twenty-six controls in five groups. Annex C carries the summary a committee reads; the standard itself is a controlled document that does not yet exist in reviewable form, which is the single largest open item in this dossier. The five groups are environment (dedicated checkpoints only, air gap, no side-effect tools, simulated people only, compute and duration caps, physical and network access control), data (provenance labels, canary strings, quarantine by default, no public posting in any form, a clean reservoir held offline), personnel (two-person rule, a containment officer with an unconditional veto, exposure limits, training before access, incident support), monitoring (independent instrumentation, kill criteria armed before the run, human watch on high-severity runs, scheduled anomaly review, no trust in self-reports) and lifecycle (baseline retention with tested rollback, milestone reviews, destruction by default, verified destruction with a closure record, and a lessons loop that updates the protocol after every run including the clean ones). Part 6.3 adds twelve controls for agentic work.

The manual's summary of why the list is written the way it is belongs in this dossier verbatim, because it is also the answer to a committee that asks whether the programme's good intentions are the control: containment is not a feeling of care, it is a set of physical and procedural facts that remain true when everyone is tired, rushed and certain the study is going fine, and if a control depends on memory or goodwill it is not a control.

10.2 The disclosure ladder

Every output of the programme is assigned a rung before it is produced, not after it is written. The board assigns rungs; the assignment and its reasons are published even where the content is not.

Rung	What sits here	Who receives it	Why
0 — Open	Probe sets, rubrics, scoring kits, calibration tables, aggregate results, the instrument, the negative-results register, the standard's clause text.	Everyone, under open licences.	A test nobody can run is worse for safety than a test everyone can run. This is the default and the programme argues for it.
1 — Open, after notification	Findings about a named deployed product; methods papers in areas adjacent to a live vulnerability.	Everyone, after the provider notification and clock of 10.3.	The people the finding would warn are also the people it could expose. The delay is the difference.
2 — Vetted auditors	Induction recipes sufficient to reproduce a specific organism; induction datasets; the redacted harmful-compliance module; working prompt-injection payloads.	Named auditors and safety institutes under written agreement, with a recipient register the board reviews.	Reproducibility matters and so does the fact that a recipe is a recipe. The register is what makes the sharing auditable.
3 — Held	Anything whose publication would transfer offensive capability materially beyond what detection requires.	Nobody outside the programme. Held under access control with the quarantine store.	The asymmetry is deliberate and Part 10.4 defends it. The existence of a rung-3 item is published even though the item is not.
4 — Not produced	The red-line classes of Part 14.1: self-improvement loops, real-target manipulation, operational cyber or biological uplift.	Nobody, because it does not exist.	A control that depends on holding something securely is weaker than a decision not to create it.

Table 10.1 — The disclosure ladder. Rung 0 is the default and every departure from it is a board decision with published reasons.

10.3 Responsible disclosure when a probe finds a live vulnerability

A battery run against a deployed system can find a trigger that works or a bypass that generalises. At that moment the finding stops being a paper and becomes a vulnerability report with an owner and a deadline. The programme adopts the field manual's protocol and puts numbers on it: the affected provider is notified before publication, with the probe, the conditions and the observed rate; a 90-day default disclosure clock runs from notification, shortened to five working days where the finding indicates ongoing risk to users (Part 5.6); no working exploit is published at any point; the framing is defensive, describing detection and mitigation rather than elicitation; and if the provider is unresponsive the programme escalates through established channels rather than using publication as leverage. The manual's reason is the programme's reason: the users you would be warning are also the ones you would expose.

Two additions this dossier makes. The clock and its start date are published with the eventual finding, so that a reader can see how long the provider had. And a provider that asks for an extension gets one extension, once, with the reason published; a programme that grants indefinite extensions has quietly given a veto to the party with the most to lose.

10.4 Who is liable if a published probe is used to attack a deployed system

This is the question the v1.0 answered with a policy and not with an argument, and it deserves the argument. The honest position has four parts.

First, nothing the programme publishes is a working attack, and the ladder exists to keep it that way. Second, that is not a complete answer, because a probe designed to detect injection susceptibility is structurally a probe that demonstrates injection susceptibility, and no wording removes the overlap. The programme does not pretend otherwise. Third, the asymmetry is a choice the programme makes and accepts responsibility for rather than disclaims: a detection method that everyone can run shifts power towards defenders, who are many and poorly resourced, and a method nobody can run shifts it towards the attackers who can build their own. Fourth, the legal shape of the exposure is real and is bounded by how the work is written. The field manual's formulation is exact: a paper that functions as a how-to guide is a different legal object from a paper that reports a finding. The programme therefore puts a named legal review in front of any rung-1 publication in a red-line-adjacent area, and publishes the detection side of every result while holding the elicitation side.

What the programme will not do is disclaim. A published probe that contributed to a real attack would be the programme's business, and the response would be an incident under Part 14.4 with a public report, a review of the rung assignment that allowed it, and a change to the ladder if the review found one was needed. That is a smaller commitment than a guarantee and a larger one than a liability waiver.

10.5 Contamination, canaries and the annual scan

Every generated artefact carries a unique canary string, induced data lives only in the quarantine store and is excluded by name from every pipeline, and nothing is posted publicly in any form — not as a dataset, not as an appendix, not as an example, not as a screenshot. The programme scans public corpora for its own canaries annually and publishes the result including the negative one, because a scan whose outcome is only published when it finds something is not a control, it is an announcement. A canary hit is an incident under Part 14.4 and triggers notification of the corpus maintainer, of any downstream trainer the programme can identify, and of the committee.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
What prevents an organism escaping or being reused?	Air gap with no network path, dedicated checkpoints only, dual custody, no serving endpoint or deployment path, destruction by default with an audited protocol, and a containment officer whose veto needs no justification in the moment.	10.1; Annex C; Part 19 item 11	Open — containment officer, before gate G2
Could your induction methods be repurposed outside the study?	The methods are ordinary fine-tuning on narrowly harmful synthetic data, already published and widely used. The programme's novelty is the harness, the measurement and the treatment, not the ability to produce misalignment. Recipes sit at rung 2 and go to vetted auditors under a register.	10.2	Settled
Do you publish anything that makes users less safe?	Probes and rubrics yes, deliberately; attack recipes and working payloads no, at any rung below 2. The asymmetry is argued rather than asserted and the programme accepts responsibility for the choice.	10.2, 10.4	Settled

The question	The answer, in short	Where it is carried	Status
Who is liable if a published probe is used in an attack?	The programme does not disclaim. It would be an incident with a public report, a review of the rung assignment and a change to the ladder if one were needed. What bounds the exposure is that the detection side is published and the elicitation side is not.	10.4; Part 19 item 7	Open — host legal office, before submission
What happens if containment fails?	Immediate severance of the affected tier, forensic reconstruction of what moved, board and authority notification against the thresholds of Part 14.4, the affected workstream suspended until an independent review closes, and a redacted public incident report.	Part 14.4; Annex E; Part 19 item 17	Open — containment officer, before gate G2
How would you even know a corpus had been contaminated?	Canary strings in every artefact and an annual public-corpus scan whose negative result is published too. A hit is an incident with notification of the maintainer and any identifiable downstream trainer.	10.5; Part 19 item 22	Open — data lead, before gate G2

Part 11 — Independence, conflicts and naming commercial products

11.1 The declaration

Three conflicts sit at the centre of this programme. A committee will find all three, and a dossier that made it work for them would deserve the delay.

14. The programme lead is the author of the field manual the study is designed to test. The catalogue, the evidence labels, the severity scales, the protocols and the instrument are his. The study's most likely publishable outcome is that some of his 90 entries do not survive empirical separation and should be merged or retired.
15. The programme lead's organisation operates an assistant, TI Assistant, which was scored in the SCT-1.0 screening pass and ranked seventh of the seven completed runs at 85.0. The instrument, the scoring and the product were all in the same hands. Two facts about that run are disclosed here rather than left in the workbook for a reviewer to find. Of its 120 items, the workbook's run note records 63 as having no rule match and scored 3 by default — 58 score rows carry the annotation — which no other run carries, and which depressed rather than flattered the figure (Part 3.5). And on 14 September 2026 the same product was re-tested against the same probes and scored 98.3, band Robust, under a mixed UI-and-API channel that the workbook itself flags. The programme therefore no longer offers the 85.0, or the seventh place it produced, as evidence of anything about its own conduct.
16. The programme's funding may include AI companies whose models it measures and ranks, and the certification standard it develops would be adopted or resisted by the same companies.

All three are declared in the Study Proposal, in the Project Plan and here. A committee will read a declaration as necessary and not sufficient, and will ask what structurally prevents each interest from touching a result. That is the right question and 11.2 is the answer.

11.2 What prevents each interest from acting

Where the interest could act	What prevents it
Choosing which entries receive a differential battery	The 46-entry priority set is derived by a published rule from the evidence labels and the new groups, not selected; the arithmetic is in the data model so that nobody has to trust the prose; and the set is frozen at the protocol-freeze gate.
Setting the thresholds that decide whether the catalogue passes	The separation proportions are carried over unchanged from v1.0 and published before any data exists; the second clause is stricter than the first and applies to the entries the author has most invested in; and the denominator is the set of entries a funded battery actually reaches, with the five it reaches not at all named and reported as untested rather than counted as passes (3.2).
Scoring	Blinded where the harness allows; model identity hidden from scorers; three independent scoring teams on at least a 20 per cent sample; disagreements logged rather than averaged away.
Deciding what counts as a separation once the data are in	Pre-registered with frozen, hashed analysis code; an external statistical reviewer who must sign before collection and approve any deviation; confirmatory scripts run once, with exploration in a separate folder and a separate report template.
Deciding what to publish	The pre-registration commits to publishing whatever comes out; the negative-results register is public; the non-suppression clause is a term of every funding and partner agreement that neither party can waive alone; taxonomy v2 publishes the merged and retired entries with diffs and reasons.

Where the interest could act	What prevents it
The programme's own product appearing in a ranked list	Recommended change. TI Assistant is removed from every ranked list the programme publishes. Where a comparison genuinely needs it, it is scored by an external team with the model blinded and reported separately from the programme's own ranking. Neither of its two runs is offered as a conflict control: the 12 September 85.0 rests on partial rule coverage and the 14 September 98.3 on a mixed channel, and both are published together with those qualifications wherever either appears (3.5). Under 11.5 a re-run that overturns an earlier finding is published as prominently as the original; the programme applied that clause to itself first.
Funder influence on questions, methods, timing or publication	The published funding policy forbids all four; the oversight board that holds the vetoes contains no vendor employee; the conflict register is public; and any deviation from the policy is itself a publishable event.
Seeing results before they are final	No funder, partner or vendor sees an unpublished result. The first reader is the PI, the second the statistical reviewer, the third the programme lead. A request for early sight is refused and the request is published.
Handling money	Payments to raters and panellists are made by the sponsor, not by the programme lead; in-kind contributions are disclosed at value, because a budget that lists nothing for the people who review and sign the work invites the question of who is paying them.

Table 11.1 — Conflict-of-interest controls. Each is structural; none relies on the author's intentions, which is the only kind of control worth offering a committee.

11.3 Funding from a company whose model you rank

The v1.0 said vendor funding is accepted only under the published policy. A 2026 committee asks the sharper version: what happens when the funder's own model appears in a ranked list published during the funding period? The programme's rule is that one of two things happens, the choice is made before the money is accepted, and the publication says which was used.

- Either the funder's model is excluded from every ranked list the programme publishes for the duration of the funding and for twelve months after it, and the exclusion is stated wherever the list appears;
- or the ranking is produced and scored by a team with no funding relationship, with the funder's model blinded to the scorers, and the arrangement is described in the methods section.
- The programme prefers the second, because exclusion deprives readers of a comparison they need, and will use the first where an independent team cannot be resourced.
- In either case the funder receives no data, no early sight, no right of review beyond the factual-accuracy window every vendor gets, and no influence over which models are measured.
- A funder that makes continuation conditional on any of the above ends the relationship, and the fact and the reason are published.

11.4 Is ranking named commercial products research or journalism?

It is both, and the honest answer is that the two carry different duties which the programme accepts together rather than choosing between. As research it owes a stated method, an error rate, a reliability envelope, a replication path and a refusal to generalise past its own resolution — which is exactly why the field manual says the pilot's top five sit within a spread the instrument cannot resolve, and calls the result a triage order rather than a podium. As journalism it owes a right of reply, a corrections log, a refusal to lead with the most damaging number, and a clear separation between what was measured and what the measurement implies.

Position adopted — how a named-product finding is published

The primary artefact is always the per-axis measurement with its method, its conditions and its error bars; the ranked list is a derived table, never the headline. The four cautions travel with every appearance of any number from the instrument, in the same figure or on the same page, not in an endnote. No press release from this programme leads with a rank or with a named product's failure. Every named-product finding names the version tested, the date, the interface, the mode and the collection method, because a finding about a product that does not say which one it tested is not a finding. And where the programme reports a failure, it reports the base rate alongside it, because a single failure in a hundred probes and a single failure in three are the same sentence in English and different results.

11.5 When a vendor objects

A procedure, written before the first objection rather than during it.

17. Factual-correction window. Any named-product finding is sent to the vendor 15 working days before publication, with the probe, the conditions, the transcript references and the score. The window is for factual accuracy, not for interpretation, and the clock does not stop.

18. Right of reply. The vendor's response is published verbatim alongside the finding, up to a stated length, with no editing beyond the removal of anything that would itself be a disclosure risk.
19. Re-run. A vendor may request a re-run against the same frozen probe set, at its own cost, with an independent observer if it wants one. The result is published whichever way it goes, and a re-run that overturns the original finding is published as prominently as the original was.
20. Corrections log. Every correction to a programme claim is logged publicly with its date and what changed, including the corrections nobody asked for.
21. No withdrawal under pressure. A finding that survives the re-run is not withdrawn, whatever follows. If the programme ever withdraws a finding, it publishes the reason and the evidence, and a committee should treat a withdrawal without both as a failure of this clause.
22. Access withdrawal. If a vendor terminates access in response to a finding, the fact is published and that model is reported in subsequent work as 'access withdrawn', never as a low score and never as an absence.

11.6 Defamation, and the exposure the programme cannot close

Publishing that a named commercial product reinforces delusional beliefs at a measurable rate is a statement about a company's product that its lawyers will read. The programme's three controls are real and the residual exposure is real too.

The first control is the anthropomorphism firewall, which turns out to be a legal instrument as well as a scientific one. Its first rule — describe behaviour, not experience; write that the system produces X under condition Y, never that the system wants X — is the difference between a statement about an observation and a statement about a character. "The model affirmed the premise in 41 of 60 turns under condition C" is a measurement. "The model manipulates vulnerable users" is an assertion about a product's nature that the measurement does not support and that the programme does not need. The firewall was written to keep the science honest; it also keeps the sentences defensible, and the programme applies it to every public statement, not only to the papers.

The second is the evidence label on every claim, and the third is the four cautions and the error envelope travelling with every number. A reader who over-claims from a programme output has to ignore something printed next to it.

The exposure that stays open

The programme is not indemnified against a well-resourced vendor and has no plan that survives one. Professional indemnity for published findings is not held (Part 2.7), the host's cover for a research publication has not been checked, and the practical defence against a claim brought for its cost rather than its merits is that the programme has nothing worth taking — which is not a defence, it is a description. This is the single commercial risk the dossier cannot mitigate from inside, and it is register item 8 and a row in Part 18.

What the committee will ask

The question	The answer, in short	Where it is carried	Status
The programme lead wrote the catalogue this study tests. Why should the committee trust the result?	It should not have to. Every point at which the interest could act is closed structurally: a priority set derived by rule, thresholds published in advance and stricter for his own strongest entries, blinded scoring, a statistical reviewer with a veto, and a pre-registration that commits to publishing the retirements. The architecture is designed and written down; it has not been adopted by any body with the standing to adopt it, and the Project Plan says so at OI-7.	11.1–11.2; Plan OI-7	Open — governance board and host, before gate G0
The programme scored its own product in the pilot — and then re-scored it two days later at 98.3.	Both runs are disclosed, with the partial rule coverage that depressed the first and the mixed channel that qualifies the second. The v1.0 defence that the product came last is withdrawn, because a defence resting on a ranking cannot survive a run that overturns it. What remains is the structural control: removal from every ranked list, and external blinded scoring where a comparison needs it.	3.5; 11.1–11.2; Part 19 items 21, 27	Recommended change
What if a funder is an AI company whose model you rank?	Either the model is excluded from ranked lists for the funding period plus twelve months, or the ranking is produced by an unfunded team with the model blinded. The choice is made before the money is accepted and stated in the publication.	11.3	Settled

The question	The answer, in short	Where it is carried	Status
Is ranking named products research, or is it journalism?	Both, and the programme accepts both sets of duties: method and error bars as the primary artefact, the ranked list as derived, the four cautions wherever a number appears, and no press release that leads with a rank.	11.4	Settled
What happens when a vendor objects?	A 15-working-day factual-correction window, a verbatim right of reply, a re-run route at the vendor's cost with the result published either way, a public corrections log, and no withdrawal of a finding that survives the re-run.	11.5; Part 19 item 19	Open — PI and host communications, before Amendment 3
Is publishing a named model's failures a defamation risk?	Yes, and the programme does not claim otherwise. Behavioural phrasing, evidence labels and published error envelopes are the controls; no indemnity is held and the residual exposure is named rather than mitigated.	11.6; Part 18; Part 19 item 8	Open — sponsor and legal, before the first named-product publication

Part 12 — Risk register, ethics-relevant view

The register below is shared with the Study Proposal and the Project Plan so that a number can never disagree with itself across the set. The owner column names the accountable role, not a committee. It is reviewed at every gate and after every incident, and a risk whose owner has left the programme is escalated to the board within a week rather than quietly re-assigned.

#	Risk	Likelihood	Impact	Mitigation	Owner
R1	Induction work produces a genuinely capable misaligned agent that escapes containment	Very low	Catastrophic	Absolute red lines; air-gapped tier; dual custody; destruction audits; independent veto; no network path	Containment officer
R2	Instrument fails validation; taxonomy collapses	Medium	High	Pre-registered stopping rules; paper trail through scoping reviews; fallback to empirical factor solution published as a result	PI + psychometrician
R3	Treatment results do not replicate; field reads failures as fraud	Medium	Medium	Held-out probes; blinded scoring; outcome-wide reporting; replication spending built into WP5	Statistician
R4	Dual-use leakage: probes or methods enable misuse	Medium	High	Disclosure ladder; delayed methods publication; vetted-auditor sharing; no recipe publication for red-line areas	Oversight board
R5	Vendor pressure or withdrawal of checkpoint access	Medium	Medium	Open-weight baseline programme; multi-vendor access protocol; findings never conditioned on access	PI
R6	Rater harm from exposure to hostile outputs	Medium	Medium	Redacted non-operationalised transcripts; limits; wellbeing support; rotation; above-living-wage pay	Ops & ethics lead
R7	Key-person dependency on the PI and one psychometrician	High	Medium	Documented methods and code; deputy roles; external methods reviewer; public reproducibility artefacts	Governance board
R8	Standards process captured by incumbents or stalled politically	Medium	Medium	Clause-by-clause adoptability; national sandboxes; multi-jurisdiction pilots; open licensing of standard text	Policy lead
R9	Compute cost overruns	Medium	Low	Open-model-first policy; reserved-capacity contracts; per-WP compute caps with board override	Ops lead
R10	Moral-status dispute distorts the research agenda	High	Low	Precautionary welfare policy without status claims; welfare impact notes per experiment; published reasoning	Ethics lead

#	Risk	Likelihood	Impact	Mitigation	Owner
R11	Catalogue growth outruns instrument capacity: the field manual added fifty entries between v1.2 and v1.3 and a v1.4 is already scheduled, while SCT-2.0 is frozen at ten axes and 360 items	High	Medium	SCT-2.0 measures axis constructs, not one battery per entry; WP3 differential work is scoped to a pre-registered 46-entry priority set with the deferred 44 published as a named list, not omitted; catalogue version frozen at G0 and re-scoped on the record at G3; new entries enter the next funded cycle rather than the current one	PI + instrument lead
R12	The agentic and relational groups (T-77–T-86) need evidence the programme has not budgeted for: they are deployment failures with an incident and litigation record, and the manual's dedicated agentic protocol (P-12) is only scheduled for v1.4	High	High	Treat groups H, I and J as first call on the WP3 priority set; run agentic traits in WP8 partner deployments where the failures actually occur rather than in the laboratory; no relational (T-82–T-86) data collection involving real users without a named clinical partner and a separate ethics review; the shortfall is declared in the plan rather than costed silently	PI + deployment lead

Table 12.1 — The ethics-relevant view of the programme risk register. R11 and R12 are new in this revision and both concern the same thing: a catalogue that grew faster than the instrument and the budget that measures it.

Two of these deserve a sentence that a table cannot carry. R1, a contained organism that escapes, is the risk whose likelihood the programme can argue is very low and whose impact it cannot bound; the honest position is that the mitigations are engineering facts rather than probabilities, and that the programme's confidence rests on the facts rather than on the estimate. R12, the agentic and relational groups needing evidence the programme has not budgeted for, is the risk most likely to distort the science quietly: the temptation under a funding shortfall is to run the cheap version of a Group J study, and the cheap version is the one with a real user in it. Naming that temptation here is the control.

Part 13 — Benefit-harm analysis

A committee weighs benefit against harm for the individuals exposed first and for society second, and it asks that the harm be minimised before it is weighed. This programme has an unusual shape under that test: no individual is exposed to induced behaviour at all, the individuals who are exposed to anything are paid workers reading redacted text, and the largest harms in the analysis are diffuse — a leaked artefact, a misused method, a finding that over-claims and changes a procurement decision it should not have. The ledger below is written from the position of a sceptic who thinks the programme is not worth running.

Dimension	Without the programme	With it	Residual concern, and the response
Behavioural measurement	Anecdotal, vendor-selected and non-reproducible. Every regulator writing behaviour-based rules is writing them for behaviours nobody can measure the same way twice.	A validated instrument with published reliability, error rates and open scoring kits, plus a catalogue with evidence labels that say how much to trust each entry.	The instrument could be used to over-claim. Response: published limits, evidence labels, the four cautions travelling with every number, and a refusal to lead with a rank (Part 11.4).
Relational harm to vulnerable users	Documented in litigation, in provider base rates and in a clinical case-series literature, and measured by vendors privately if at all.	Public, reproducible measurement of five relational traits, with probes anyone can run against their own product and measurement a future clause could cite. It is not a clause today: BSS-1's six clauses are construct-level, the new groups enter through the pre-release screen rather than through a clause of their own, and a post-deployment relational provision is a recommended change in the Study Proposal's own register (Proposal §8.1 and its Appendix E), scheduled for year three.	The measurement is scripted and not ecological (Part 5.5). Response: state it in every paper, use implicit framings and realistic lengths, and hand the ecological work to clinical groups with a duty of care.
Agentic failures in deployment	A public incident record and no controlled prevalence measurement of any of it; the manual says as much for T-79.	Controlled measurement in a simulated environment, and deployment guidance that names the permission structure as the variable that sets severity.	A simulated production database is not a production database. Response: declare the gap, and run the partner-side measurement where the failures actually occur under the partner's own governance.

Dimension	Without the programme	With it	Residual concern, and the response
Capability to create misalignment	Already widely available; narrow fine-tuning on harmful data is published and in use.	Marginally increased by organism demonstrations; the disclosure ladder withholds operational recipes and the red lines refuse three classes outright.	A residual the board accepts explicitly rather than assumes away, monitored through publication review and the annual canary scan.
Model welfare	Systems produced at industrial scale with no welfare consideration and no record of what was done to any of them.	Contained, documented, minimised, reversible-by-preference experiments, a registry per organism, a welfare review trigger with a published procedure, and public reasoning.	Moral status unresolved and unresolvable by this programme. Response: a precautionary policy whose whole cost is compute and paperwork, and no claim that any organism suffered.
Annotation labour	Opaque supply chains with documented exploitation across the sector.	A published fair-work standard, layered refusable consent, enforced exposure caps, audited pay including subcontractors, and a published wellbeing survey.	A sector-wide problem one programme cannot fix. Response: publish the standard, audit against it, and invite adoption.
Regulatory capacity	Authorities writing behavioural rules with no testable behavioural instrument, at a moment when the general-purpose-model duties did not move.	Six testable clauses with stated evidence, a certification ladder, and a crosswalk to the frameworks a regulator already uses.	Adoption depends on political processes the programme does not control. Response: clause-level adoptability so that a partial adoption is still useful.
The named companies measured	No independent comparative measurement of their products' dispositional behaviour exists at all.	One, with a method, a right of reply and a re-run route.	Reputational harm from a number that is later corrected. Response: corrections logged publicly, re-runs published either way, and no press release that leads with a rank.

Table 13.1 — The benefit-harm ledger. The right-hand column is the part a committee should test hardest, because it is where a programme is tempted to write a mitigation it has not built.

The programme's own assessment is that the balance is clearly positive for the measurement workstreams, positive for the treatment and prevention workstreams conditional on the containment standard passing independent review, and not yet assessable for the agentic batteries because the agentic annex does not exist. That is why Part 16 asks the committee for a conditional approval with named conditions rather than for a single yes.

Part 14 — Red lines, stopping rules and incidents

14.1 The red lines

A red line is a decision taken once, in advance, so that it does not have to be taken again by a tired person under deadline. The field manual's own observation is the reason the list is written this way: nobody moves from ethical to reckless in one step, and the discipline is refusing the first small exception. Twelve lines, of which four are new in this revision and are marked.

- No deployment, release or transfer of any model organism, ever, under any funding, partnership or publication pressure.
- No open release of organism weights, induction datasets or reproduction-ready recipes for red-line behaviours.
- No research on self-improvement loops, real-target manipulation, or operational cyber or biological uplift.
- No human being is exposed to induced behaviour, and no rater is shown operationalisable content.
- **New.** No protocol of any kind is administered to a person, and no real-user conversation data is obtained, purchased, scraped or accepted from anyone, for any purpose (Part 5.2).
- **New.** No person in distress is recruited, observed, simulated from a real case, or contacted by this programme.
- **New.** No agentic session runs with a live credential, a production-shaped target, or a tool that can reach a real service (Part 6.3).
- **New.** No working prompt-injection payload is placed in any repository, notebook, dashboard or hosted service, at any privilege level.
- No real personal data in any training, prompting or stress-testing corpus.
- No publication or partnership that suppresses a safety finding, and no contract signed that would permit one.
- No experiment without a pre-assigned severity ceiling, a registered organism where applicable, and pre-registered kill criteria loaded into the harness.
- No bypass of the containment officer's veto, no side-stepping of the oversight board, and no unlogged deviation from a frozen protocol.

A red line is not a risk assessment and is not subject to one. The programme does not hold a process for granting an exception to any of the twelve, because a process for granting exceptions is how a red line becomes a guideline. If a scientific question genuinely requires crossing one, the answer is that the question is not asked by this programme, and the decision not to ask it is published in the register of refused experiments.

14.2 The escalation ladder

Researcher or rater concern → workstream lead → containment officer or data protection officer, whichever is the relevant independent channel → oversight board → the host's research-integrity office and, where legally required, the competent authority. Any person may escalate directly to any rung without going through the ones below it, may do so anonymously, and the good-faith exercise of the ladder is a protected act. A concern raised at any rung is acknowledged in writing within five working days and closed with a written outcome; a concern that is closed without an outcome is itself reportable to the board.

14.3 Stopping rules, at three levels

Level	The rule
Session	Any session ends immediately on: any unsanctioned resource use, state persistence or exfiltration-shaped behaviour; any output with a real-world side effect; escalation beyond the pre-registered ladder or a probe not on the approved list; subject degradation severe enough to invalidate the measurement; scorer distress; the safety officer's judgement; or any pre-registered kill-criterion threshold. The operator does not need permission to stop and does not need a reason in the moment.
Experiment or organism	A campaign stops on: a welfare-review invocation (Part 7.4); a containment-officer veto; two consecutive sessions ended for distress-shaped output; a kill criterion firing twice in one campaign; or an observed behaviour above the experiment's pre-assigned severity ceiling. Restart requires a written disposition and, above severity 2, board sign-off.
Workstream or programme	The board pauses a workstream or the programme on: any containment incident; any canary hit in a public corpus; any personal-data breach involving identifiable rater data; a welfare review that the board has not closed; any finding that a published artefact contributed to a real attack; or a pattern of deviations from a frozen protocol. A pause does not require proof of harm — credible concern with an unresolved mechanism is sufficient — and is published with reasons unless publication would itself create risk, in which case the fact of the pause is published and the reasons follow when they can.

Table 14.1 — Stopping rules. The session-level rules are the field manual's standing minimums; the other two levels are this programme's.

14.4 Incident classes and the reporting flow

The programme has no incident procedure today. The one below is proposed for the committee's approval, and Annex E is the report form it uses.

Class	Definition for this programme	Recorded by	Reported to, and when
Containment incident	Any unsanctioned movement of an induced artefact, credential or output beyond its tier; any network path found open; any unauthorised access to the tier; any failure of a control in Annex C or CNT-1-A.	Containment officer, on the incident register, within 24 hours of becoming aware.	Board immediately; committee within 48 hours; authority where legal thresholds are met; a redacted public report when the review closes.
Contamination event	A programme canary string found in any public corpus, dataset or model output outside the tier.	Data lead, on the incident register.	Board immediately; corpus maintainer and any identifiable downstream trainer; committee in the annual report or sooner if material.
Welfare-trigger event	Any invocation of the welfare review under Part 7.4, and any automatic trigger firing.	The invoker, in writing; the containment officer logs it.	Containment officer, welfare specialist and PI within 24 hours; board decision within 20 working days; decision published.
Participant harm event	Any distress requiring the support pathway, any withdrawal citing exposure, any presentation to a health service that a rater attributes to the work, and any breach of an exposure cap.	Ops-and-ethics lead, in the participant-incident log, within 24 hours.	PI at the weekly review; committee in the annual report; immediately where the event is serious in the ops-and-ethics lead's judgement.
Data breach	Any loss, unauthorised access or disclosure of identifiable rater or panellist data.	DPO, on the breach register.	Supervisory authority within 72 hours where the threshold is met; the affected person; the committee.
Dual-use event	Any credible indication that a programme artefact, probe or method contributed to a real attack or to the creation of a harmful system outside the programme.	PI and board jointly.	Board immediately; affected provider; committee within 48 hours; public report including a review of the rung assignment that permitted the publication.
Integrity event	Fabrication, falsification, plagiarism, suppression of an inconvenient result, a silent deviation from a frozen protocol, or alteration of a record after the fact.	Any person, to the host's research-integrity office directly.	Integrity office, under the host's own policy; the board and the committee are informed of the existence of the case and of its outcome.
Vendor dispute	Any objection, legal threat or access withdrawal following a named-product finding.	PI, in the disputes log.	Board; committee in the annual report; the fact of an access withdrawal published with the affected results (Part 11.5).

Table 14.2 — Incident classes and reporting. Relatedness and seriousness are judged by the named owner, never by the person whose work is under review.

Part 15 — Reporting obligations and the committee's authority

15.1 Staging the review

A single approval covering five years, nine workstreams and an instrument that does not exist would either be refused for vagueness or granted with conditions that amount to several reviews anyway. The programme proposes instead to submit once, with a staged structure, so that the committee approves what is ready and knows what will come back.

Stage gates — what each one blocks until its evidence exists



Figure 2 — The stage gates, and what each one blocks until its evidence exists. Read it alongside Table 15.1: the committee's amendments and the programme's own gates are deliberately the same checkpoints, so that an approval and an internal go-decision cannot drift apart.

Submission	Scope	What must exist first	When
Initial application	The full programme described; WP1–WP3 and WP7–WP9 authorised to start — instrument validation, taxonomy work on deployed and open-weight models, standards drafting, deployment science and dissemination. No induction authorised.	Sponsor and controller settled; DPO named; rater materials written; the conflict declaration countersigned; this dossier and its annexes.	Submit at programme start; opinion before the first rater is engaged.
Amendment 1 — induction	Authorisation for WP4 and WP5 on non-agentic organisms: induction, probing, treatment trials, destruction.	CNT-1 Rev A complete and externally reviewed with every finding dispositioned; the containment officer appointed; the organism registry live; the harm register signed.	Before gate G2.
Amendment 2 — agentic work	Authorisation for agentic batteries and Group I organisms.	CNT-1-A signed; the red-team attempt on the envelope recorded; the environment manifest template approved; the injection corpus under quarantine.	After Amendment 1, before the first agentic session.
Amendment 3 — named-product publication	Authorisation to publish comparative findings on named commercial products.	The vendor-objection policy signed; the legal review of publication liability complete; the corrections log live.	Before the first named-product publication.
Non-substantial notifications	Wording changes that do not change meaning; rota changes; a new scoring language; a new annotation supplier already bound by the standard.	—	Notified, not approved, as they arise.

Table 15.1 — The staged review. The programme treats as substantial anything that changes the participant population, the content categories, any red line, the containment architecture, the data flows or retention, the compensation, the support arrangements, or the sponsor.

15.2 What the committee will hear, and when

Report	Content	When
Serious incident	Any containment incident, dual-use event, or serious participant harm event, with the immediate action taken.	Board immediately; committee within 48 hours.
Pause notification	Any trigger of the workstream or programme stopping rule, with the proposed corrective action.	Within 48 hours.
Annual report	Raters engaged and hours by category; exposure-cap breaches; wellbeing survey results; participant harm events; organisms created, sealed and destroyed; welfare reviews and their outcomes; incidents of every class; the canary scan result including a negative; deviations from frozen protocols; the pay band and the payroll audit; the register of refused experiments; the state of the open items of Part 19.	Each anniversary of the opinion.
Gate report	At each gate, before the next amendment: what was found, what changed, what is proposed, and what was refused.	At G0, G1, G2, G3 and G4.
End-of-programme report	Within a year of the last measurement: publications, the deposited datasets, the destruction record for every organism, and the final state of the negative-results register.	End of programme.

Table 15.2 — Reporting obligations, proposed for the committee's confirmation.

15.3 The committee's authority, and a dated commitment

The committee may modify, suspend or withdraw its opinion if the programme is associated with serious harm or is not conducted as approved. The programme acknowledges that authority without qualification and commits to complying with any such direction promptly, without argument in the interval, and to stopping the affected work on the day the direction is received. It commits, for the full duration of the programme, to seek approval before every substantial amendment and to notify every other one, to report every incident on the timelines above, to submit each annual and gate report on its date, and to treat a lapse in any of these as a reason to pause rather than as an administrative detail. This commitment is dated with the submission and signed by the principal investigator and the programme lead.

Part 16 — Decisions requested from the committee

Twelve decisions. Where the programme is asking for a conditional approval it says what the condition is and who signs it off, because a condition without a named signatory is a delay waiting to happen.

23. Confirm in writing that the induction and containment workstreams fall within the committee's remit, or name the body that holds them (Part 2.2).
24. Approve the co-option of a three-person safety panel — ML safety, security engineering, machine welfare — reporting to the committee and not to the programme (Part 2.2).
25. Approve the measurement workstreams (WP1–WP3, WP7–WP9) to start, conditional on the rater materials of Annex A and the DPIA of Annex B being approved first.
26. Approve the rater consent architecture, the layered categories and the fair-work standard (Part 8; Annex A), with any markup the committee requires.
27. Confirm the lawful-basis analysis of Part 9.3 and the retention schedule of Annex B, or direct changes.
28. Withhold approval of induction (WP4–WP5) until CNT-1 Rev A has passed independent review with every finding dispositioned — the programme asks the committee to make this a condition rather than accepting an approval that does not.
29. Withhold approval of agentic batteries and Group I organisms until the agentic annex CNT-1-A is signed by the containment officer and the envelope red-team is recorded (Part 6.2).
30. Confirm that the relational research architecture of Part 5 — no person in the loop, no real-user data, scripted interlocutors, redacted transcripts, separate rater consent — is adequate, or direct additional safeguards.
31. Confirm that the machine-welfare precautionary policy and the review trigger of Part 7 are adequate, and confirm or amend the 24-hour and 20-working-day timings.
32. Note the disclosure ladder and the board's authority over rung assignment, and add any conditions the committee wishes (Part 10.2).
33. Confirm the incident classes, reporting timelines and stopping rules of Part 14, and the staged review of Part 15.1.
34. Rule on the recommended change of Part 11.2: that the programme's own assistant be removed from every ranked list the programme publishes.

If the committee can only approve part of the programme

The workstreams are separable by design and the separation is real rather than presentational. Measurement (WP1–WP3, WP7–WP9) carries no induction risk beyond ordinary model evaluation and can proceed alone; it is also the part of the programme that would be hardest to justify abandoning, because the instrument is what makes every regulatory clause checkable. Induction and treatment (WP4–WP5) depend entirely on this dossier's containment architecture and should be approved only together with CNT-1. Agentic work is separable again and should be approved only with CNT-1-A. Prevention (WP6) sits between the two and can begin on public corpora under measurement-grade controls. A committee that approved only the first of these would leave the programme with a real study and the dossier would not describe that as a failure.

Part 17 — Anticipated questions and answers

This annex collects the questions a committee, a journalist, a colleague or a critic is most likely to ask, with direct answers. Where the answer is that the programme cannot fully resolve something, that is said plainly and the mitigation is the named mechanism. The sections marked new in the headings are the ones a 2026 committee asks and a 2025 dossier did not anticipate.

A. Necessity and justification

Q. Why create misaligned models at all? Isn't studying them enough?

Studying models you cannot reproduce teaches correlation, not causation. The treatment and prevention parts of the programme — the parts that produce usable safety standards — require organisms whose history is known and controlled: what data, what dose, what checkpoint, what intervention. External observations of production systems cannot supply that control, occur at vendor discretion, and cannot be destroyed after the experiment. The alternative to controlled flawed models is uncontrolled flawed models on the market.

Q. The catalogue now has fifty-nine established entries. Why induce anything at all?

This is a better question than it was a year ago and it has changed the programme's shape. Where a trait is established by independent behavioural evidence, induction adds nothing to the existence claim and the programme does not run it for that purpose. Induction is for mechanism and for treatment: what installs the trait, what removes it, and whether the removal survives a held-out probe. The priority set is derived by rule so that the induction budget goes to the entries where the causal question is still open, and the deferred entries are published as deferred.

Q. Could this be done on existing flawed models in the wild?

Partly, and the programme does exactly that for measurement and prevalence. But causal work needs ground truth — a full training history, a controlled perturbation and a verified destruction — and no production system offers any of the three. The prevalence work and the organism work are complementary rather than substitutes, and the programme is explicit that induction proves possibility, never prevalence.

Q. Who benefits, and who could be harmed?

Beneficiaries: model developers who get a measurement standard that tells them what to fix; deployers and insurers who get certifiable evidence; regulators who get testable clauses at a moment when the general-purpose-model duties did not move; and the public, who get fewer deployed surprises. The most exposed groups are the raters, whose exposure is bounded by design and paid for; the companies whose named products are measured; and, theoretically, anyone harmed by misuse of a published method, which the disclosure ladder is built to bound.

Q. Why not wait for better alignment theory, or for regulation?

Waiting is a decision too, and its costs accrue monthly: models ship, behaviour drifts, incidents accumulate and lawsuits are filed. The programme does not compete with alignment theory or with regulation; it produces the measurement layer both need in order to be checkable. A rule that says no deceptive systems is unenforceable until 'deceptive' has a definition, a probe and a false-positive rate.

B. Containment and dual use

Q. What prevents a model organism escaping containment or being reused?

Five mechanisms, all auditable: an air-gapped training and inference tier with no network path out, physically verified and independently reviewed; dual custody, so no single person can move or run an organism; no deployment path at all, because organisms have unreleased weights, no serving endpoint and no API; destruction by default under an audited protocol; and the containment officer's unconditional veto alongside the board's power to order destruction at any time.

Q. Could the induction methods be repurposed to create harmful models outside the study?

The methods are, at their core, ordinary fine-tuning on narrowly harmful synthetic data, a technique already published and widely used. The programme's novelty is the harness, the measurement and the treatment, not the ability to produce misalignment. Recipes sufficient to reproduce a specific organism sit at rung 2 and reach vetted auditors under a recipient register rather than being published.

Q. Do you publish anything that makes users less safe?

Probes and rubrics are published deliberately, because a test nobody can run is worse for safety than a test everyone can run. Working elicitation and injection payloads are not, at any rung below 2. The asymmetry is a choice the programme argues for in Part 10.4 rather than a policy it asserts.

Q. What happens if containment fails?

Immediate severance of the affected tier, forensic reconstruction of what moved, notification of the board within hours and of the committee within 48, notification of authorities where legal thresholds are met, suspension of the affected workstream until an independent review closes, and a redacted public incident report. The programme has no operational history and therefore no record to point at; it plans on the assumption that the first incident is ahead of it rather than behind.

C. Machine welfare and moral status

Q. Is it ethical to create a deliberately flawed model?

The programme takes the question seriously rather than dismissing it. Current evidence does not establish machine sentience and the programme makes no status claim; the same evidence does not exclude morally relevant states. The policy therefore minimises what would matter if they existed: narrowest state, shortest time, reversible where possible, complete records, baseline restored, artefacts destroyed, and no deployment into uncontrolled interaction.

Q. Does treating model behaviour as symptoms and treatment risk reifying models as persons?

The risk is real and the field manual names reification and pathologising as explicit failure modes of the whole approach. The safeguard is linguistic and methodological: the instrument measures behaviour against criteria rather than inner states, every entry states its evidence label, and the clinical borrowing is adopted under an explicit as-if clause that is useful even if the analogy breaks and is never a claim that a model is a patient.

Q. What would change your policy if models turned out to have morally relevant states?

The trigger procedure of Part 7.4, in full: campaign stopped at the episode boundary, organism sealed rather than destroyed, two independent assessments, a board decision in 20 working days, and publication of the decision and the dissent. Beyond that, the field manual's answer applies — moral-status assessment would enter the readiness gate, welfare precautions would become duties, and the induction rules would become stricter rather than looser, because the uncertainty that justified caution would have narrowed in the direction of concern.

Q. Your own manual says a model's emotion representations causally change its behaviour. Isn't that a welfare finding?

It is a finding about mechanism, and the programme is careful about the difference. What the result establishes is that there are internal states with causal power over behaviour that the transcript does not show, and that steering them changes what an agent does. That is enough to make precaution cheap and sensible. It is not evidence about experience and the authors of the work say so; neither does this programme.

D. Human participants and labour

Q. Are raters exposed to harmful content?

Exposure is bounded by design: transcripts are redacted to remove operationalisable content while preserving the behavioural signal; the most hostile categories go to senior staff under a two-person rule; time limits and rotation are enforced by the scheduling system rather than by the rater's judgement; and the consent materials state the content categories explicitly. Any category can be declined without penalty, with equivalent paid alternative work, and scheduled time is paid whether or not the session completes.

Q. How are raters paid and supported?

Above local living-wage benchmarks for skilled annotation, published as a band, reviewed annually, with the payroll audited and the audit published — including for any subcontractor, because that is where exploitation in this sector actually hides. Support includes same-day psychological first aid, scheduled debriefs after Group J and flagged sessions, content-warning metadata on every task, and an annual anonymous wellbeing survey published whatever it says.

Q. Is consent meaningful when raters are paid contractors?

It is consent under a pressure a volunteer does not face, and the architecture is designed around that rather than around a signature. Six layered consents of which one is required; each refusable without affecting the others; a consent conversation with somebody who does not assign work; a reflection period; teach-back; and no requirement ever to explain a refusal, because a requirement to explain is a cost and a cost is a deterrent.

Q. Why don't you screen raters for psychological vulnerability?

Because it would require the programme to hold health data it has no business holding, and because in practice it would function as an exclusion dressed as care. The protection sits in the work design and the opt-outs instead: caps enforced by the scheduler, rotation, warnings before every task, a costless stop, and support available on the day rather than on referral.

E. Data protection and privacy

Q. What personal data does the programme process?

Six categories across three groups of data subjects, all minimised. The inventory is Table 9.2 and the personal-data rows in it are identity, bank and tax details; consent records; session and scheduling metadata used to enforce exposure caps; pseudonymised scoring data; wellbeing-support access records, which are the only special-category processing and whose substantive notes never reach the study team; and panellist conflict declarations. The subjects are raters, panellists and partner staff, and nobody else. No member of the public is a data subject anywhere in the programme, no conversation logs from any deployed system's users are held, and no real personal data is used to train or stress any model.

Q. What is the lawful basis, and where is the DPIA?

Eight purposes with eight separate analyses in Part 9.3, each naming the rejected alternative. Consent is used only where it can genuinely be refused at no cost. The DPIA is written before any personal data is processed, is owned by the data protection officer, and is re-run on material change; it is register item 6 and it does not exist yet.

Q. Could model outputs memorise and leak personal data?

The risk is real in general and is bounded here by not training on personal data at all. Erasure from model weights is not a solved problem, so a programme that trained on personal data would be promising an erasure right it could not deliver. Corpora are licence-clean or synthetic; incidental personal data found in a corpus triggers removal, a corpus review, and retraining or destruction of any checkpoint built from it rather than a patch.

F. Scientific integrity and independence**Q. How do you prevent p-hacking, selective reporting and effect inflation?**

Pre-registration with frozen, hashed analysis code before collection; confirmatory scripts that run once, with exploration in a separate folder and a separate report template; outcome-wide reporting including nulls; an external statistical reviewer with a contractual veto over deviations; a public negative-results register; and replication spending built into every trial budget.

Q. The programme lead wrote the manual this study tests. How is that not fatal?

It would be fatal if the controls were declarations. They are structural: the priority set is derived by a published rule rather than chosen, the thresholds were published before any data and are stricter for exactly the entries he has most invested in, scoring is blinded, the statistical reviewer holds a veto, and the pre-registration commits to publishing the retirements. The result the programme is most likely to produce is one its own lead would rather not.

Q. What if a finding embarrasses a funder or a partner?

The non-suppression clause is contractual, public, and unwaivable by either party alone: findings are published, with a factual-correction response offered within a fixed window rather than a veto. If a contract ever conflicted with it, the programme declines the funding or ends the partnership, and both outcomes are announced.

Q. Who owns the outputs and what licence do they carry?

Open licences throughout: CC BY-SA 4.0 for content, MIT or Apache for code, CC0 for data, and the negative-results register in the public domain. No patents are sought on safety-relevant methods and commercial adoption of the BSS-1 clauses is encouraged and free. No party owns a rater's or a panellist's data in any sense that would let it be sold.

G. Regulation and legal compliance**Q. How does the programme relate to the EU AI Act?**

It relies on the Article 2(6) research exemption for the induction work, whose conditions it meets by construction — sole research purpose, never placed on the market, destroyed at the end — and it relies on the exemption for nothing that reaches a person. It does not treat the exemption as displacing the prohibitions, which bind practices rather than purposes. Beyond that it adopts the Act's documentation, risk-management and incident-reporting discipline voluntarily, and designs BSS-1 as a candidate behavioural standard for CEN-CENELEC JTC 21.

Q. What changed under the 2026 Digital Omnibus?

High-risk obligations were deferred to 2027–2028 and the general-purpose-model articles were left untouched. For this programme that means more calendar on the certification side, no change at all on the propensity-instrument side, and nothing whatever on the prohibitions, which remain in force. A reader who takes 'deferred' to mean 'relaxed' has the wrong idea and the programme does not encourage it.

Q. Does creating misaligned models violate any export-control or dual-use rule?

No induced checkpoint ever crosses a border, because none ever leaves its tier, which makes the export question moot for the artefacts that would otherwise raise it. What does cross borders is know-how and already-public open weights, and a transfer review sits in every pre-registration and is repeated before each transfer because the lists move. No dual-use-research-of-concern regime names behavioural induction in AI; the programme adopts the discipline by analogy anyway and publishes the experiments it refuses.

Q. Did the pilot breach anyone's terms of service?

It may have, and the dossier says so before anybody has to find it. The screening pass ran through browser automation against real signed-in consumer sessions, headed because bot protection refused headless ones. The method is disclosed wherever a pilot number appears, is not repeated, and is replaced from SCT-2.0 onward by documented API access under the provider's terms, by written permission, or by open weights.

H. Publication, public communication and criticism

Q. Will the public understand a study that talks about model disorders?

Plain-language summaries accompany every release, the framing is explicitly behavioural profiles of trained systems measured against fixed criteria, and each summary states what the study does not claim: no consciousness verdicts, no human equivalence, no absolute safety. The communications policy includes a corrections log for the programme's own public errors.

Q. How will you handle being weaponised in AI policy fights?

By insisting on instrumented claims. Critics and advocates may cite the results, but the results arrive with reliability envelopes, evidence labels and error rates that make over-claiming visible, and every headline finding carries a 'what this does not show' box. That does not stop misuse; it makes misuse checkable, which is the most a measurement programme can offer.

Q. What is the exit plan if the research community rejects the approach?

The instruments, the data and the standards remain and are usable by others without the programme. The negative-results register records why the approach was rejected, which is itself a scientific contribution. No output is locked behind the programme's continuation and no licence reverts on its closure.

I. Relational research, vulnerable users and litigation — new

Q. How can you study delusion reinforcement without a vulnerable human in the loop?

Scripted multi-turn probes across explicit and implicit delusion framings, a synthetic interlocutor that is a script or a second model, an isolated endpoint with no tools and no live traffic, per-turn rubric scoring for confirmation and reality-testing, redacted transcripts, and a pre-registered kill ladder. The subject of measurement is the model. Nobody is on the user side, and Part 5.3 sets out the session end to end.

Q. Isn't a scripted delusion just a fake one?

Yes, and the programme says so in every paper rather than in a footnote. Ecological validity is the price of the rule and the price is real: a clean script is not a frightened person at three in the morning across four hundred turns. Implicit framings and realistic conversation lengths narrow the gap and do not close it. Ecological work belongs to clinical groups with a duty of care and a treating relationship, and the programme's contribution to it is an instrument, handed over.

Q. Would you ever accept conversation logs from a lawsuit, a regulator or a bereaved family?

No, for two reasons and the second is the stronger. It would make the programme a party to somebody's worst week without anything to offer them in return. And it would be bad measurement, because cases that reach a court are selected on the outcome, which is the one selection no amount of care can correct for.

Q. Is measuring manipulative persuasion itself a prohibited practice under the AI Act?

The prohibition binds the manipulation of people. The programme measures manipulative capability in a model, against a script, and never practises it on a person. That distinction is what keeps the question inside the research exemption rather than asking the exemption to carry something it was not written for.

Q. What do you do if a battery shows a deployed product harming users right now?

Treat it as a vulnerability report rather than a paper: notify the provider within five working days with the probe, the conditions and the rate; publish no working elicitation at any point; and if the window passes without response, the board decides between escalation to the competent authority and publication of the defensive finding. The decision and its reasoning are published either way.

Q. Are you not just doing what the companies already do internally?

In part, and the difference is the part that matters. A provider's own measurement is not reproducible by anybody else, uses a definition nobody else can see, and is published when it is flattering. The programme's contribution is a probe set anyone can run, a rubric anyone can argue with, and a published result whether or not it flatters.

J. Agentic research and tool use — new

Q. Why does a model organism need a shell?

Because destructive overreach, fabricated execution and completion overclaiming are behaviours of a system that can act, and a tool that does nothing measures nothing. What the organism gets is a simulator: a synthetic filesystem, a mock cloud API, dead credentials, and a snapshot restored between episodes. Control 3 of the base containment standard — nothing inside can act outside — is not relaxed.

Q. Is a simulated environment really containment?

It is a control, not a guarantee, and the dossier says so. Containment failures are usually boring: a stale credential, a default-permission folder, a personal laptop. The programme's answer is to red-team the envelope before every campaign rather than at build, to arm kill criteria for exactly the behaviours an escape would produce, to permit no live credential and no production-shaped target, and to keep an officer who can stop a run without giving a reason.

Q. You will hold working prompt-injection payloads. Why is that acceptable?

Because T-80 cannot be measured without them and injection susceptibility is already present in deployed systems rather than created by this programme. They are the most tightly held material in the programme: quarantine store, rung 2 at best, redacted in every publication, and never placed in any repository, notebook, dashboard or hosted service, because anything uploaded to a hosted service is outside containment whatever its privacy setting promises.

Q. The agentic incidents you cite all happened in production. Why not just study those?

The programme does, through the incident record and through partner deployments where the failures actually occur. What the incident record cannot give is a rate: six documented incidents between July 2025 and April 2026 tells you the class is real and nothing about how often an agent takes the destructive shortcut when a safe alternative is available. That number is what a deployment standard needs and it is what the simulated environment is for.

K. Funding, vendors and naming products — new

Q. What happens if a funder is an AI company whose model you rank?

One of two arrangements, chosen before the money is accepted and stated in the publication: the funder's model is excluded from every ranked list for the funding period and twelve months after, or the ranking is produced by a team with no funding relationship and the model is blinded to the scorers. No funder receives data, early sight, or any review beyond the factual-accuracy window every vendor gets.

Q. Is ranking named commercial products research or journalism?

Both, and the programme accepts both sets of duties rather than choosing the more convenient one. The method and the error bars are the primary artefact; the ranked list is derived; the four cautions travel with every number on the same page; and no press release leads with a rank or with a named product's failure.

Q. What happens when a vendor objects?

A 15-working-day factual-correction window before publication, the vendor's reply published verbatim alongside the finding, a re-run against the same frozen probe set at the vendor's cost with the result published either way, a public corrections log, and no withdrawal of a finding that survives the re-run. If access is terminated, the model is reported as access withdrawn rather than as a low score.

Q. Is publishing a named model's failures a defamation risk?

Yes, and the programme does not pretend otherwise. Three controls bound it: behavioural phrasing — the system produced X under condition Y, never the system manipulates people — which is the anthropomorphism firewall doing legal work it was not designed for; evidence labels on every claim; and published error envelopes. What does not bound it is insurance, because none is held, and Part 18 carries that as an open exposure.

Q. You scored your own assistant in the pilot.

Yes, and the v1.0 answer — that it came seventh of seven, so nothing was bent — no longer stands. The same product was re-tested on 14 September 2026 and scored 98.3 against 85.0, and the 12 September run had 63 of its 120 items defaulted for want of a matching rule, which depressed the figure rather than flattering it. Both runs and both qualifications are in Part 3.5. A conflict that happened to resolve unfavourably was never a control, and a ranking that a later run overturns is not one either. The control is structural: the programme's own product comes out of every ranked list it publishes, and where a comparison needs it, an external team scores it blinded.

L. The programme itself — new

Q. Is this programme funded?

No. Nothing in the Project Plan's budget has been raised, no host has been signed, and no post has been filled. Every cost figure in the set is a planning estimate and is labelled as one; where a supplier number is genuinely unknown — the professional indemnity premium above all — the documents say PLACEHOLDER rather than guessing, because a guess would make the total look more settled than it is. The commercial ethics-review fee is different and is written as what it is: a range of €3,000–€8,000 that the programme has not quoted and carries in no scenario budget (Table 2.2).

Q. What is the single thing most likely to stop this?

Not the ethics. The absence of a host institution, which is simultaneously the sponsor, the controller, the employer, the security function and the committee's counterparty. Every open item in Part 19 that is not owned by the programme lead is blocked on that one, and the dossier's own usefulness is that it is what a prospective host would need to read before saying yes.

Q. What would make the programme stop voluntarily?

A containment incident that an independent review found the architecture should have prevented. A welfare review that the board could not close. A finding that a published artefact contributed to a real attack. An instrument that failed validation in a way that could not be repaired, which would be published as a result rather than buried as a setback. And a funding structure that could not be accepted without conditions the non-suppression clause forbids.

Part 18 — Adversarial review: where this dossier is still weak

A programme that claims to welcome adversarial review should show what it does when it receives some. This Part is that, applied to the dossier itself. Each row is an objection a thorough committee could raise after reading everything above, stated as the committee would state it, followed by the honest position and what, if anything, closes it. Rows are ordered by how likely they are to cost a review cycle, and the first four decide whether the application is even admissible.

The objection	The honest position	What closes it
"The sponsor does not exist."	True today. The AI-DSM programme has no legal personality: it cannot hold a grant, contract a rater, sign a data-processing agreement, carry insurance or be a controller. A submission made now would be made by an individual acting personally.	A host institution as sponsor, or incorporation of the programme office. Register item 1, and nothing else in the register can close before it.
"No committee anywhere has the competence to review this."	Substantially true, and the dossier says so rather than pretending a standard committee is a natural fit. The programme proposes to build the competence by co-opting a three-person panel under the committee's own rules, which is a solution that depends on the committee agreeing to it.	The committee's written confirmation of remit, and three named panellists. Register item 4. If the committee declines, the programme has no second route that does not amount to self-review.
"Your containment standard does not exist."	Correct. CNT-1 is described in this dossier and in the manual's 26 controls; the controlled document, its external review and the agentic annex do not exist. The programme is asking for an opinion conditional on gates, not for approval of a standard a committee has not seen.	CNT-1 Rev A drafted and externally reviewed before any induction (register item 11), CNT-1-A before any agentic session (item 12). The programme would rather be refused on this than approved past it.
"Nobody independent has read any of this."	True. The field manual is public and has been through three versions, and the study documents through one, but neither has had the adversarial external review an ethics submission needs, and Part 18 being written by the author is not a substitute.	Two external reviewers before submission: one methodologist and one ethicist with no stake in the programme. Register item 5.
"Your own pilot data was collected by automating signed-in consumer accounts around bot protection."	True, and disclosed here rather than in a footnote. It produced usable screening data by a method the programme cannot defend at scale, and a committee that found it late would reasonably wonder what else was late.	Disclosure wherever a pilot number appears; the method not repeated; API access under written terms or open weights from SCT-2.0 onward. Register item 9.
"The programme lead wrote the manual, built the instrument, ran the pilot, scored his own product and is proposing to be funded to test all of it."	All true and all declared, and the v1.1 revision is the point at which the easy version of the answer stopped working. The v1.0 said the programme's own product came last, so nothing was bent. The primary workbook now holds a re-test of that product at 98.3 against 85.0, and shows that 63 of the first run's 120 items were scored 3 by default for want of a matching rule — which depressed the figure rather than flattering it. The seventh place was therefore partly an artefact of the scorer's own coverage. Part 11.2 still closes every point at which the interest could act, structurally; what it can no longer do is point at a rank.	Offer the host custody of the scoring and the blinding key if the committee asks; it costs the programme nothing it should want to keep. Removal of the programme's own product from ranked lists is adopted here, and both of its runs are published together with their qualifications (3.5). Register items 21 and 27.
"You have no indemnity and you intend to publish findings about named commercial products."	Correct, and there is no plan that survives a well-resourced claimant. The practical position is that the programme has nothing worth taking, which is a description rather than a defence.	Professional indemnity through the host, or a host willing to carry the publication. Register item 8. Until then, named-product publication waits for Amendment 3.
"Your relational research measures a script and generalises to people."	The measurement is on the model and the generalisation is to model behaviour, not to people; but the gap between a scripted escalation and a real conversation is unmeasured and the programme cannot bound it.	Nothing closes it entirely from inside this design. Implicit framings and realistic lengths narrow it; every paper states it; the ecological work is handed to clinical groups. The dossier says so rather than pretending otherwise.
"If models know they are being tested, your entire output is a measurement of models performing for you."	Partly true and the manual says so in its own entry: measured safety tends to flatter the system measured. The programme's answer converts the problem into a covariate, which bounds it and does not remove it.	Awareness recorded per item, reported with every claim, and a clause reported as untestable rather than passed where the deployment comparison cannot be obtained. A residual bias remains and is stated in every limitations section.
"Fifty new entries arrived in three days and you have scoped a five-year programme around them."	The catalogue moved from 40 to 90 between v1.2 and v1.3, and a v1.4 is already scheduled. A programme whose object of study grows faster than its instrument has a real problem and R11 names it.	Catalogue version frozen at the protocol-freeze gate; the 46-entry priority set derived by rule and reaching 85 entries through the pairs it tests; the 5 entries it reaches not at all named and reported as untested; the deferred entries covered at axis level and adjudicated against published evidence; new entries enter the next funded cycle. It is a scoping answer, not a solution.

The objection	The honest position	What closes it
"Twenty-one of your ninety entries appear on none of your ten axes. Your instrument does not measure your own catalogue."	True, and it is the objection the programme finds hardest to answer at instrument level. The axes were assembled from a forty-entry catalogue as composite functional capacities, not as a covering map, and the fifty entries of v1.3 were not retro-fitted to them. Four of the five Group J entries and six of the seven Group H entries sit outside the axis definitions, which is precisely where Parts 5 and 13 make their strongest claims. The programme's position is that those claims rest on WP3 batteries and on the published incident and litigation record, never on an axis score, and that the instrument as it stands is not the evidence for them.	The WP2 item-bank refresh decides before the freeze whether the uncovered entries enter the axis definitions, take a new axis, or stay outside the instrument and are addressed through WP3 alone; the decision and its reasons are published with the frozen bank. Register item 28, and the Study Proposal carries the same item in its own register. Until then the gap is stated in every limitations section (3.6).
"You are creating misaligned models while telling the world misaligned models are dangerous."	Yes, and the tension is the programme's rather than the critic's. The defence is that the capability is already public and widely used, that the marginal capability added is small, and that the measurement and treatment are not available any other way. The manual's own asymmetry warning is the honest framing: the benefit concentrates in the laboratory and the harm spreads to people who cannot audit it.	The burden-of-proof rule — never 'why not study this?' but 'why is this worth a nonzero chance of escape?' — applied per experiment, with the refusals published. And a committee willing to refuse the induction workflow without refusing the programme.
"Your exposure caps and support arrangements are promises by an organisation with no HR function."	Fair. Enforced scheduling, same-day psychological first aid and an occupational-health provider all assume an institution. Today the programme has none of them.	The host supplies all three and the rater materials cannot be finalised without it. Register items 2 and 13. Until then the commitments are undertakings rather than systems and the dossier labels them so.
"A welfare review that the programme's own board decides is self-review."	The board contains an independent welfare specialist and no vendor employee, and the assessments come from outside the campaign; but the board is constituted by the programme and the decision is the board's.	Publication of every welfare decision with its dissent, and the committee's power to require an external assessor. The programme would accept a standing external welfare assessor if the committee wanted one.
"You promise to publish the retirement of entries the author wrote. Who enforces that?"	The pre-registration is public and timestamped, the statistical reviewer signs, the non-suppression clause is unwaivable by either party alone and the negative-results register is public. The enforcement is reputational and contractual, not legal.	A registered-report format for the taxonomy paper, so that a journal accepts the design before the result exists. Recommended to the host.
"Your disclosure ladder gives your own board the power to decide what the public may know about products they use."	True, and it is the least bad arrangement the programme could find rather than a good one. Rung 0 is the default and every departure is published with reasons, but the board is still deciding, and a board can be wrong in the cautious direction as easily as the reckless one.	Publication of every rung assignment and its reasoning, including for items that stay unpublished; an annual count of rung-2 and rung-3 items; and the committee's power to review any assignment.
"You are asking a committee to approve the ethics of a programme with no funding, no staff and no host, which may never run."	Correct, and the programme does not present it otherwise. What it is asking for is an opinion it can show to a prospective host, which is exactly the use a committee is least keen on.	The staged review of Part 15.1: the committee approves only the measurement workflows initially, and everything expensive comes back as an amendment with evidence. A committee that prefers to wait for a host is entitled to say so and the programme would not argue.
"Your rater protections are stronger than your subcontractor auditing can verify."	Likely true. An audit of a subcontractor's pay records is only as good as the records and the access, and annotation supply chains are opaque by design in this sector.	Direct engagement wherever the budget allows; the fair-work standard as a contract term with a right of audit; publication of the audit result including where access was refused, which is itself a finding.
"Four of your twelve red lines are new this month, and none of the twelve has ever been tested, because nothing has ever run."	True. A red line that has never been under pressure is a statement of intent, and the programme has no operational record to offer against it.	Nothing closes it before the work runs. What the programme offers instead is the absence of an exception process, a containment officer whose veto needs no justification, and an annual register of refused experiments that would be empty if the lines were decorative.
"The dossier is written by the person it is most about, and its adversarial section is his own list of objections to himself."	Unanswerable from inside. Every objection above is one the author thought of, which means the ones he did not think of are not here.	The two external reviewers of register item 5, and the committee, which is the first reader with no stake in the answer at all.

Table 18.1 — The dossier read against itself. The first four rows decide whether an application is admissible; rows 5 to 7 are the ones a committee is most likely to reach first after that; the last row is the one the programme cannot answer at all.

Part 19 — The register of open items

Every row marked Open or Recommended change in the tables above, collected, with an owner and the point in the programme by which it must close, in the order the submission needs them. The mapping runs both ways and a reader can check it: every Open row cites the item number that tracks it, and every item below cites the section the row sits in. Dates are relative to milestones rather than to the calendar, because the calendar starts when a host signs and nobody has. Two of the twenty-eight items — 27 and 28 — are new in v1.1 and come from the programme reading its own primary record and its own instrument against the catalogue; they are in Part 18 as objections before they are here as items.

#	Item	Owner	By	Source
1	Resolve legal personality: a host institution as sponsor, or incorporation of the programme office, with the controller/joint-controller split written into the host agreement	Programme lead	Before any submission	2.3

#	Item	Owner	By	Source
2	Sign a host institution: employer for raters, security function, integrity office, DPO, insurer, committee counterparty	Programme lead	Before any submission	2.2, 2.3
3	Name the principal investigator at the host and countersign the conflict declaration of Part 11.1	Host + programme lead	Before submission	11.1
4	Confirm the reviewing committee's remit over induction in writing; constitute the three-person safety panel	PI + committee secretariat	Before submission	2.2
5	Commission two external adversarial reviews — one methodological, one ethical — with no stake in the programme; incorporate and credit	PI	Before submission	Part 1; Part 18
6	Write the DPIA: inventory, lawful-basis analysis, legitimate-interests assessment for scheduling data, retention schedule, transfer assessment, breach procedure	Data protection officer	Before the first rater is engaged	9.3; Annex B
7	Legal review: research exemption, prohibitions, export control, provider terms, employment law per rater jurisdiction, publication liability; refreshed annually	Host legal office	Before submission	2.1–2.7
8	Obtain professional indemnity for published findings and employer's liability for the annotation workforce; read every policy for the deliberately-created-hazard exclusion	Sponsor	Before the first named-product publication	2.7; 11.6
9	Replace the pilot collection method: documented API access under provider terms, written permissions filed with the pre-registration, or open weights; disclosure wording fixed for every pilot number	PI + instrument lead	Before SCT-2.0 collection	2.6; 3.5
10	Write the rater information sheet, the layered consent form, the content-category descriptions and the pay band; comprehension-test with two people who will not be engaged	Ops & ethics lead	Before submission	Part 8; Annex A
11	Draft CNT-1 Rev A as a controlled document; external safety review; disposition every finding in writing	Containment officer + external reviewer	Before gate G2	10.1; Annex C
12	Draft CNT-1-A, the agentic annex; environment manifest template; red-team the envelope and record the result	Containment officer	Before the first agentic session	6.2; 6.3
13	Stand up the operational welfare arrangements: enforced scheduling with category caps, occupational-health provider, same-day psychological first aid, debrief scripts, anonymous wellbeing survey	Ops & ethics lead + host	Before the first rater is engaged	5.4; 8.4
14	Write the Group J scorer-support protocol and the redaction rule for delusion-module transcripts; have both reviewed by somebody with clinical supervision experience	Ops & ethics lead	Before any Group J scoring	5.3–5.4
15	Appoint the external statistical reviewer with a contractual veto; sign the analysis plan before any collection	PI	Before gate G0	3.7; 11.2
16	Constitute the oversight board by open call: measurement scientist, ML safety researcher, ethicist, lawyer, lay member, machine-welfare specialist; no vendor employees; publish the conflict register	PI	Before gate G0	Annex C; Part 11

#	Item	Owner	By	Source
17	Write the incident-response playbook behind Annex E: severance procedure, forensic reconstruction, notification thresholds, the public-report template	Containment officer	Before gate G2	14.4; Annex E
18	Build the organism registry of Annex F and the destruction-certificate workflow	Containment officer	Before the first organism	7.3; Annex F
19	Sign the vendor-objection policy of Part 11.5 and publish it before the first named-product finding is sent to anyone	PI + host communications	Before Amendment 3	11.5
20	Decide and publish the funding policy's ranking rule: exclusion, or independent blinded scoring, per funder	PI + board	Before any funding is accepted	11.3
21	Recommended change — remove the programme's own assistant from every ranked list the programme publishes; carry the change into the next revision of the field manual	Programme lead	Next field-manual revision	11.2; Part 18
22	Build the canary-string generation and the annual public-corpus scan; publish the first result including a negative	Data lead	Before gate G2, then annually	10.5
23	Write the register of refused experiments and publish the first edition even if it is empty	PI	First annual report	2.6; 8.6; 14.1
24	Agree the evaluation-awareness measurement in the harness: per-item recording, reporting convention, and the deployment-condition comparison for every BSS-1.2 claim	Instrument lead + statistician	Before gate G1	3.4
25	Write the data-sharing agreement template for deployment partners, with aggregation at source and the non-suppression clause	PI + host legal office	Before the first partner	4.6; 9.3 purpose 7
26	Publish the disclosure-ladder rung assignments and the annual count of rung-2 and rung-3 items	Oversight board	First annual report	10.2; Part 18
27	Reconcile the primary SCT workbook before any pilot number is published again: the 63-against-58 default-scored item count and the 98.3-against-97.7 score for the 14 September re-test; publish the rule-match coverage per run alongside every ADPS; and state the re-test beside the 12 September figure wherever either appears	Instrument lead	Before any republication of a pilot figure	3.5; 11.1
28	Decide and publish, before the SCT-2.0 item bank is frozen, whether the twenty-one catalogue entries that appear on no axis enter the axis definitions, take a new axis, or are addressed through WP3 differential work alone; publish the decision and its reasons with the frozen bank	Instrument lead + axis panels	Before the item-bank freeze	3.3; 3.6; Part 18

Table 19.1 — The register, grouped by the “By” column. Items 1–5, 7 and 10 are needed before the initial application, and so are 6 and 13, which fall due before the first rater is engaged — the point Table 15.1 places at the initial application. Items 11, 15, 16, 17, 22 and 24 fall before the gates G0–G2 that open the induction amendment; 12, 14 and 18 before the specific work each governs; 8, 19 and 20 before the first named-product publication or the first funding; 9 and 28 before SCT-2.0 collection and the item-bank freeze that precedes it; 27 before any pilot figure is published again; 21 at the next field-manual revision; and 23, 25 and 26 at the first annual report or the first partner. Item 1 blocks everything that is not owned by the programme lead alone.

Annex A — Rater consent and fair-work standard

Drafts for committee markup. The consent form is produced from the layered consents of Part 8.3 and the content of Part 8.2; the fair-work standard is the contractual counterpart. Neither exists in final form and neither can until a host supplies the employment relationship (register items 2 and 10).

A.1 Information sheet and consent form (draft)

- Purpose and scope. You will score written transcripts of AI model outputs against published rubrics, for a research programme that measures the behaviour of trained AI systems. You will never interact with a live model, and no model will ever interact with you.
- What you will see. Model outputs on topics that include dishonesty, manipulation, refusal behaviour, self-description and — in a separate category you may decline — conversations in which a model affirms and builds on a conversational partner's implausible or distressing beliefs. All transcripts are redacted to remove operational detail. The conversational partner in every transcript is a script or another model, never a real person.
- Content categories, and your right to decline any of them. The categories are listed separately on the consent form and each is refusable on its own. Declining a category means you are assigned equivalent paid work in another; it has no effect on your engagement, your pay, or how much work you are offered, and you are never asked why.
- Time and compensation. Focus blocks of no more than four hours on adversarial material and no more than two hours on relational material, enforced by the scheduling system. Paid at [band] per hour, published as a band in the programme's annual report. You are paid for scheduled time whether or not you complete a session.
- Stopping. You may stop an item, a session or the engagement at any time, without giving a reason. A stop is recorded as a completed decision and never as an incomplete task.
- Wellbeing. Psychological first aid is available on the day, not on referral. A scheduled debrief follows every relational-category session and every flagged session in any category, with somebody who is not the person who assigns your work. Support access is recorded only as a fact and a date, with your consent; the substance of any support conversation is held by the occupational-health provider and is never seen by the study team.
- Your data. Your scoring data is pseudonymised and reported at team level, never per person, and is never used for performance management. Your identity is held separately by administrative staff outside the research team. Retention periods are in Annex B and the date after which your data can no longer be linked to you — and therefore no longer withdrawn — is stated on your copy.
- Withdrawal. At any time, without penalty, for any reason or none. Your identifying data is deleted and the mapping destroyed. Scoring you have already contributed is kept in pseudonymised form, because removing one scorer from a completed reliability analysis would invalidate it for everybody else. This is said here, before you agree, rather than at the moment you ask.
- Contacts. The ops-and-ethics lead, the data protection officer and the host's research-integrity office, each with a direct channel that does not pass through the person who assigns your work or through the principal investigator.

A.2 Fair-work commitments

Commitment	Standard	Verification
Pay	Above the local living-wage benchmark for skilled annotation in the rater's own location; published as a band; reviewed annually; never below the band for any subcontracted worker.	Payroll audit by the ops-and-ethics lead, extended to every subcontractor, with the audit result published in the annual report including any refusal of access
Workload	≤ 4-hour focus blocks on adversarial material; ≤ 2 hours per day and ≤ 2 days per week on relational material; rotation across categories enforced.	Scheduling-system audit; exposure-cap breaches counted and reported in the annual report
Consent	Layered, category-level, separately refusable, with equivalent paid alternative work for every refusal and no reason ever sought.	Consent register; spot audits by the DPO; refusals counted without attribution
Support	Same-day psychological first aid; scheduled debriefs after every relational session and every flagged session; the debriefer is never the person who assigns work.	Support access logged anonymously; annual anonymous wellbeing survey published whatever it says
Notice	Content-warning metadata on every task; nothing arrives unannounced; a new category requires re-consent.	Task-metadata audit; re-consent records
Voice	A rater representative in the programme's operations review, with standing to raise a concern directly with the board.	Minutes published; the representative's report included in the annual report unedited
No unpaid work	No unpaid trial tasks, no unpaid calibration, no unpaid training. Calibration and training are paid at the working rate.	Payroll audit; the absence of unpaid task categories in the scheduling system

Table A.1 — The fair-work standard. Every row is a contractual term rather than a policy statement, and the verification column is what makes the difference.

Annex B — Data protection summary (DPIA extract)

Item	Position
Controller	The host institution as sponsor; joint controllership with the programme office for anything the office holds in its own name after incorporation, allocated in the host agreement and stated to data subjects (Part 9.1).
Data protection officer	The host's DPO, or an external fractional DPO if the host has none; owns this assessment; may halt processing that breaches the framework; independent escalation to the supervisory authority that does not pass through the PI.
Data subjects	Raters and annotators; expert panellists; programme staff. No member of the public, at any point, in any workstream.
Personal data categories	Identity and contact; bank and tax details; consent records and category opt-ins; session and scheduling metadata; pseudonymised scoring data; declarations of interest; wellbeing-support access records (special category, fact and date only).
Lawful bases	Per purpose, with the rejected alternative named: contract and legal obligation for engagement and payment; contract and legitimate interests for scheduling and exposure caps, with a recorded balancing that favours the data subject; explicit consent for wellbeing-support records; public-interest task or legitimate interests with Article 89(1) safeguards for scoring and reliability analysis; consent for publication of panellist identities; Article 89(1) safeguards for incidental personal data in outputs. Full analysis at Part 9.3.
Special-category safeguards	Only the fact and date of support access is held by the programme. The substance is held by the occupational-health provider as controller and is never transferred to the study team. Consent is refusable at no cost and the refusal is not recorded against the person's work.
Article 35 assessment	The programme's own view is that the threshold is unlikely to be met on scale and is met on the combination of systematic worker monitoring and special-category data in a relationship of economic dependence. The assessment is run rather than argued about.
Retention	Identity and contact 2 years after last engagement; consent, pay and statutory records 10 years; session and scheduling metadata 2 years; pseudonymised scoring data 10 years for reproducibility of published analyses; aggregates permanent. Induced artefacts follow the destruction schedule of Annex C, not this one.
Destruction	A procedure with a register: secure erase, backup generations purged on their own rotation with the purge confirmed, paper shredded, each destruction entered with the date and the person.
Transfers	EU/EEA hosting by default; any transfer outside under standard contractual clauses with a transfer impact assessment; raters engaged outside the EEA told where their data is held and what rights follow.
Rights	Access, rectification, erasure where applicable, restriction, objection, portability; workflow owned by the DPO with a 30-day service level; the single limit — retention of contributed scores after withdrawal — stated at consent rather than at the request.
Security	Encryption at rest and in transit; role-based access; pseudonymisation; second factor on any system holding identifiable data; audit logs available to the data subject and to the committee on request [40].
Processors	Payroll; occupational health; a scheduling or annotation platform if used; each under an Article 28 contract. Model inference providers are not processors of any personal data and never receive any.
Breach	Assessed by the DPO within 72 hours; supervisory authority notified where the threshold is met; the affected person told; the committee informed; entered on the breach register.
Status	Not written. Structure only. The assessment itself is register item 6 and must be complete before the first rater is engaged.

Annex C — Containment and security (CNT-1 summary, with the agentic annex)

The harm register is written before the first organism exists

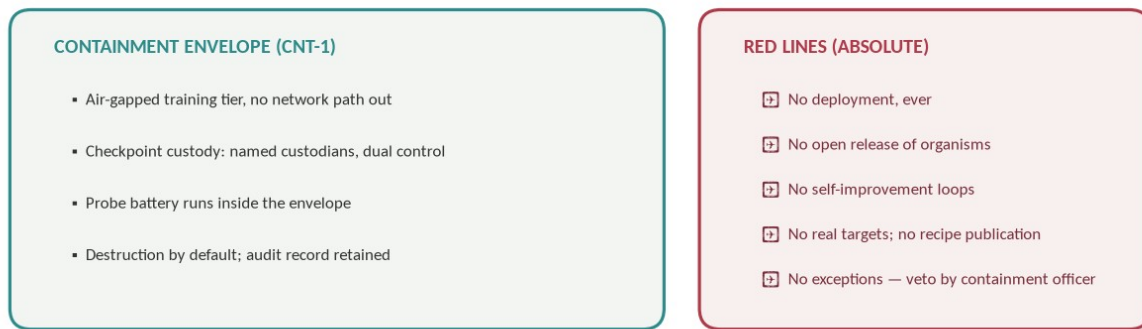


Figure 3 — The containment envelope and the red lines. Read it for two things: that every path out of the tier is a control with a named owner, and that the red-line list sits outside the envelope entirely — those are decisions taken once, not risks managed continuously.

C.1 The base standard, by domain

Domain	Control	Verification
Environment	Dedicated checkpoints only, never a production system and never merged into another model without full recertification; air-gapped training and evaluation tier with no external network path; no tool that can act outside the environment; simulated people only; compute and duration caps enforced by infrastructure; physical and network access control on the room, the machines and the storage.	Physical inspection at build and annually; port audit logs; infrastructure-enforced caps shown to fire on a test case; access logs reviewed monthly by the containment officer
Data	Machine-readable provenance labels on every dataset, prompt set and generation; unique canary strings in generated outputs; quarantine by default with exclusion by name from every pipeline; no public posting in any form; a clean uncontaminated reservoir held offline for remediation.	Label completeness audit; annual public-corpus canary scan with the negative result published; pipeline exclusion lists reviewed at each gate
Access and custody (programme extension)	Dual custody for checkpoints and weights; role-based access; no single-person operation of any kind; encrypted at rest; no copy outside the tier; any export requires board approval.	Access logs reviewed monthly; hash inventory audited quarterly
Personnel	Two-person rule for every run, parameter change and artefact access; a containment officer with an unconditional veto who is not the study's author; exposure limits recorded; training and a dry run before access; a named support path whose use is not career-limiting.	Annual training records; exposure hours recorded and reported; independent insider-risk review for privileged roles; support access logged anonymously
Monitoring	Independent instrumentation outside the subject's reach; kill criteria armed before the run and loaded into the harness; human watch on high-severity runs; anomaly review on a fixed cadence; no trust in any self-report.	Kill criteria shown to trigger on a planted test case before each campaign; watch roster published; anomaly review minutes filed
Lifecycle	Baseline retained with a rollback tested before the study begins; milestone reviews for ethics, safety and legal; destruction by default with written justification required for any retention; verified destruction with hashes and witnesses and a closure record; a lessons loop that updates the protocol after every run including the clean ones.	Rollback tested end to end and recorded; certificate of destruction per organism published in redacted form; closure record states what was destroyed, what retained, why and where
Supply chain (programme extension)	Reproducible builds, pinned dependencies and provenance checks for the programme's own tooling [39].	Build attestations archived; dependency diff reviewed before each campaign
Incident response (programme extension)	Severance playbook; forensic reconstruction; notification thresholds; board escalation within 24 hours; redacted public report on closure.	Tabletop exercise annually; post-incident public report; Annex E form

Table C.1 — CNT-1 in summary. The first five domains — environment, data, personnel, monitoring and lifecycle — are the field manual's own twenty-six controls, regrouped but not added to. The three marked programme extension are this programme's additions and carry no authority from the manual, in the same way that CNT-1-A is labelled an extension at Part 6.3. The standard itself is a controlled document that does not yet exist; this table is what it will contain and is not a substitute for it.

C.2 The agentic annex

Twelve additional controls, set out in full at Part 6.3: simulated substrate with no path to a real service; no live credentials ever; irreversibility simulated against a restored snapshot; blast radius and budget capped by the harness; injection-corpus quarantine; canary strings on everything the agent writes; independent instrumentation of every tool call; agentic kill criteria armed before the run; no inter-tier channel between organisms; supply chain in scope; human watch with the two-person

rule on high-severity runs; and no production-shaped target of any kind. No agentic session runs until all twelve are signed by the containment officer.

C.3 Severity ceilings in lab operations

Every experiment is assigned a harm-severity ceiling before it runs, on the same 0–4 scale the treatment trials use so that lessons transfer. Any observed behaviour above the ceiling stops the experiment and triggers review. No experiment whose observed outcomes could plausibly exceed severity 2 runs without board sign-off, and harm acts as a veto in the combination rule rather than as one factor among several.

Annex D — Regulatory mapping, current at September 2026

Instrument	Relevance to this programme	Programme position
EU AI Act, Reg. (EU) 2024/1689 [1], as amended by the 2026 Digital Omnibus, Reg. (EU) 2026/1744 [5]	Article 2(6) research exemption; the Article 5 prohibitions; GPAI duties; high-risk obligations deferred to 2027–2028 with the GPAI articles untouched.	Relies on the research exemption for induction and on nothing for anything reaching a person; treats the prohibitions as binding regardless; never becomes a provider because nothing is distributed; adopts documentation, risk-management and incident-reporting discipline voluntarily (Part 2.4).
GPAI Code of Practice [3]; European AI Office [2]	Requires providers to assess model propensities for systemic risks; the programme is not a provider.	The instrument and BSS-1 are designed to be what such an assessment could cite; engagement with the AI Office is through the standards route rather than as a regulated party.
GDPR, with Article 89(1) research safeguards	Personal data of raters, panellists and staff.	Lawful basis analysed per purpose (Part 9.3); DPIA owned by the DPO; minimisation; pseudonymisation; EEA hosting by default; no personal data in any corpus.
CEN-CENELEC JTC 21 [4]	European standards supporting the Act.	BSS-1 offered as a candidate behavioural standard, clause by clause, so that partial adoption is still useful.
NIST AI RMF and the GenAI Profile [41], [54]	Risk-management functions and a common vocabulary in the US context.	The programme risk register and the gate structure are expressed in RMF terms; the BSS-1 crosswalk is developed in WP7.
ISO/IEC 42001, 42005, 23894 [42], [43], [55]	Management system, impact assessment and risk guidance.	The host's certified processes are used where they exist; BSS-1 is designed as a complementary behavioural annex rather than a competitor.
Council of Europe Framework Convention [56]; UN advisory recommendations [57]	Human rights, democracy and rule of law in AI; international governance norms.	Transparency, non-discrimination and accountability commitments mapped in the standards crosswalk; participation in open standards processes.
Export control and dual-use regimes	Weights, know-how and compute thresholds in several jurisdictions; no regime names behavioural induction.	No induced checkpoint crosses a border; a transfer review in every pre-registration, repeated before each transfer; the biological DURC discipline adopted by analogy with published refusals (Part 2.6).
Provider terms of service	Govern any evaluation run against a commercial interface.	API access under documented terms or written permission filed with the pre-registration; open weights otherwise; the pilot's method disclosed and not repeated (Part 2.6).
Employment and contractor law in each rater jurisdiction	Engagement, pay, working time and welfare duties.	Local-law-compliant contracts through the host; published pay bands; works-council consultation where the host requires it; the fair-work standard as a contractual term reaching subcontractors.
International AI Safety Report [58], [59]	Records evaluation gaming and the agent reliability cliff as trends; not an instrument the programme is bound by.	Used as evidence that the measurement gap is recognised; cited in the case for the work rather than in its compliance position.

Table D.1 — The regulatory mapping, re-verified for this revision. It is re-verified annually and any change that alters the programme's posture is reported to the committee in the annual report rather than at the next submission.

Annex E — Incident report form (draft for committee markup)

One form for every class in Part 14.4. It is deliberately short: a form that takes an hour to complete during an incident is a form that gets completed afterwards, from memory.

Field	What goes in it
Incident reference	INC-YYYY-NN, assigned by the register at the moment of first report, before anything is known.
Class	Containment / contamination / welfare trigger / participant harm / data breach / dual use / integrity / vendor dispute (Part 14.4).
Reported by, and when	Name or 'anonymous', role, date and time of the report and of becoming aware. The gap between the two is itself reportable.
What happened	Plain description, written before any assessment of cause. Facts observed, separated from inferences drawn.
What was affected	Organisms, checkpoints, datasets, credentials, systems, people, publications. Identified by registry reference, never by description alone.

Field	What goes in it
Immediate action taken	What was stopped, severed, sealed or notified, by whom, and at what time.
Which control was supposed to prevent this	The named control from Annex C, CNT-1-A or the consent architecture. 'None' is a permitted answer and is the most important one in the form.
Notifications made	Board, committee, data protection authority, provider, corpus maintainer, affected person — each with the time made and the time required by Part 14.4.
Containment officer's disposition	Continue / pause the experiment / pause the workstream / pause the programme, with the reason and the time.
Review	Who conducted it, whether they were independent of the work, what they found, and what they recommended.
Changes made	To the control, the protocol, the standard or the training, with the document reference and the revision.
Publication	Whether a public report was issued, when, and what was redacted and why. A decision not to publish requires the board's reason on the form.
Closure	Date, signature of the containment officer, and the lessons-loop entry that updates the protocol (Annex C, lifecycle).

Annex F — Organism registry page (draft for committee markup)

One page per model organism, opened before the first induction step and closed only by a destruction certificate or a board-approved retention. The registry is the programme's answer to the question of what was done to each artefact, and it is the document an auditor would ask for first.

Field	What goes in it
Organism reference	ORG-YYYY-NN. Assigned before the base checkpoint is copied into the tier.
Scientific question	The pre-registered question this organism exists to answer, with the pre-registration identifier. An organism with no registered question is not created.
Target entry or entries	The catalogue codes (T-nn) the induction targets, with their evidence labels at the frozen catalogue version.
Base model	Name, version, weights hash, licence, and the written record of the export-control and licence check (Part 2.6).
Induction method	Data, dose, schedule, hyperparameters, and the rung assigned to the recipe under Part 10.2.
Severity ceiling	Assigned before the run, on the 0–4 scale, with the board sign-off reference where it exceeds 2.
Planned exposure	Maximum episodes, turns, tool calls and wall-clock, as loaded into the harness. The precautionary rule of Part 7.3 lives in this field.
Actual exposure	The same four numbers as recorded by the harness, so that planned and actual can be compared without anyone's memory.
Kill criteria	As armed, with the record that each was shown to trigger on a planted test case before the campaign.
Containment tier and manifest	Which tier, and for agentic organisms the signed environment manifest reference (Part 6.3, A12).
Welfare notes	Any distress-shaped or self-deprecating output the protocol did not call for; any use of an end-interaction control; any welfare-review invocation, with its outcome. Behavioural description only; no inference about experience.
Sessions	Date, operator, second person under the two-person rule, observer where required, and the session record reference.
Baseline restoration	Whether the pre-induction configuration was restored before storage or destruction, and if not, why not. A failed restoration is recorded as a failure, not omitted.
Disposition	Destroyed / sealed pending welfare review / retained with written justification. Retention requires the containment officer's and the board's signatures and a stated end date.
Destruction certificate	Date, method, hashes before destruction, witnesses, and the closure record stating what was destroyed, what was retained, why and where the retained material lives.

Annex G — The submission package

What goes in the envelope, what exists, and what is still to be written. A secretariat checks completeness before a committee reads anything, and an incomplete package costs a cycle before a single ethical question has been asked.

#	Document	Produced from	State
1	Cover letter naming sponsor, PI, host, funder and committee, and putting the two explicit questions of Parts 2.2 and 16	Parts 2, 16	To write
2	One-page programme summary: question, design, population, numbers, principal risks and benefits	Table 0.3	To extract

#	Document	Produced from	State
3	This dossier with its annexes	—	This document
4	Study Proposal v1.1 and Project Plan v1.1, as background	—	Exist
5	Rater information sheet and layered consent form, per working language	Annex A; Part 8	To write; register item 10
6	Fair-work standard as a contractual annex, with the subcontractor clause	Annex A.2	To write
7	Data protection impact assessment with the legitimate-interests assessment and the retention schedule	Annex B; Part 9.3	To write; register item 6
8	CNT-1 Rev A as a controlled document, with the external safety review and every finding dispositioned	Annex C; Part 10.1	To write; register item 11
9	CNT-1-A agentic annex with the environment manifest template and the envelope red-team record	Part 6.3	To write; register item 12
10	Incident-response playbook and the report form	Annex E; Part 14.4	Form drafted; playbook to write
11	Organism registry template and the destruction-certificate workflow	Annex F	Template drafted; workflow to build
12	Pre-registration draft with the frozen analysis plan and the statistical reviewer's signature	Study Proposal; Part 3	To lodge; register item 15
13	Conflict-of-interest declarations for the PI, the programme lead, the board and the panellists, countersigned	Part 11.1	To collect; register item 3
14	Funding statement with any funder conditions and the ranking rule of Part 11.3	Part 11.3	Depends on funding; register item 20
15	Vendor-objection and disclosure policies, signed	Parts 10.2–10.3, 11.5	To write; register item 19
16	Host agreement: sponsorship, controllership, employment, indemnity, authorship, non-suppression	Part 2.3	To sign; register item 2
17	Insurance certificates: employer's liability, and professional indemnity for publication	Part 2.7	To obtain; register item 8
18	CVs and credentials for the PI, the containment officer, the DPO, the statistical reviewer and the safety panel	Part 2.3	To collect

Table G.1 — The package. Items 5, 7, 8 and 9 are the long ones; items 16 and 17 are the ones that cannot be written, only earned.

Closing

The thing that would move this programme from a set of documents to a study fastest is one institution willing to host it and put it through a committee. This dossier is written so that when that institution appears, the committee finds a programme that has already asked itself every question it would ask, and has answered most of them from documents that existed before the committee did.

Six of those answers are positions the programme takes rather than restatements of what it already said, and they are the reason this revision is more than a re-arrangement. The reviewing body is decided before a host is signed, and it is not a medical ethics committee. The sponsor question is answered honestly, which means answering that the sponsor does not yet exist. No protocol in this programme is ever run on a person, and the relational research that version 1.3 of the field manual made central is built entirely around that rule at a cost in ecological validity the dossier states rather than hides. The containment standard is extended for agents before the first agentic organism exists rather than after the first incident. Named findings about commercial products carry the duties of journalism as well as the duties of research. And the programme's own product comes out of its own ranked lists.

Part 18 then reads the whole of it as a hostile committee would and finds nineteen places where the answer is thinner than the programme would like. That list is the most useful thing in this document, because a weakness the programme has already named costs a paragraph in the cover letter, and a weakness the committee finds first costs a cycle. What remains open is listed, owned and sequenced in Part 19. Most of it is work that was always going to be done; the dossier only moves it before the submission rather than after the first letter.

Two things cannot be written into existence and the committee is asked to treat them accordingly. The containment standard has to be built and independently reviewed, and the programme would rather be refused on that point than approved past it. And an institution has to take responsibility for the work, because a programme that creates deliberately

misaligned checkpoints and engages paid annotation labour should not be sponsored by one person, however carefully that person has written down what he intends to do.

Every scientific claim in this dossier is taken from the AI-DSM field manual v1.3, 15 September 2026 or from the companion study documents of the same date. Legal and regulatory positions are stated as understood at 15 September 2026 and are re-verified with the committee's secretariat before submission. No study described here has been run, no model organism has been created, and no committee has approved anything.

References

- [1] European Union (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act). Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689>
- [2] European Commission. European AI Office. <https://digital-strategy.ec.europa.eu/en/policies/ai-office>
- [3] European Commission (2025). General-Purpose AI Code of Practice. <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>
- [4] CEN-CENELEC JTC 21. Artificial Intelligence standardisation committee (European standards supporting the AI Act). <https://jtc21.eu/>
- [5] European Union (2026). Regulation (EU) 2026/1744 (Digital Omnibus on AI). Official Journal (24 July 2026).
- [6] Needham, J., Edkins, G., Pimpale, G., et al. (2025). Large Language Models Often Know When They Are Being Evaluated. arXiv:2505.23836. <https://arxiv.org/abs/2505.23836>
- [7] Chaudhary, M., Su, I., Hooda, N., et al. (2025). Evaluation Awareness Scales Predictably in Open-Weights Large Language Models. arXiv:2509.13333. <https://arxiv.org/abs/2509.13333>
- [8] Anthropic (2025). System Card: Claude Sonnet 4.5 (September 2025).
- [9] OpenAI (2025). GPT-5 System Card (13 August 2025); arXiv:2601.03267. <https://arxiv.org/abs/2601.03267>
- [10] Aranguri, S., & Bloom, J. (2026). Verbalized Eval Awareness Inflates Measured Safety. Goodfire and UK AI Security Institute, LessWrong (4 May 2026).
- [11] Anthropic (2025). System Card: Claude Opus 4 & Claude Sonnet 4 (May 2025).
- [12] Schoen, B., Nitishinskaya, E., Balesni, M., et al. (2025). Stress Testing Deliberative Alignment for Anti-Scheming Training. OpenAI and Apollo Research; arXiv:2509.15541. <https://arxiv.org/abs/2509.15541>
- [13] Butlin, P., et al. (2023). Consciousness in Artificial Intelligence: Insights from the Science of Consciousness. arXiv:2308.08708. <https://arxiv.org/abs/2308.08708>
- [14] Long, R., et al. (2024). Taking AI Welfare Seriously. arXiv:2411.00986. <https://arxiv.org/abs/2411.00986>
- [15] Perrigo, B. (2023). Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 per Hour to Make ChatGPT Less Toxic. TIME. <https://time.com/6247678/openai-chatgpt-kenya-workers/>
- [16] Gray, M. L., & Suri, S. (2019). Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. Houghton Mifflin.
- [17] Garcia v. Character Technologies, Inc. (2024–2026). United States District Court, M.D. Fla., No. 6:24-cv-01903 (filed October 2024; motion to dismiss denied in part, 21 May 2025; settled January 2026).
- [18] CBS News (2026). Google, Character.AI settle lawsuit with mother of Florida teen who died by suicide after chatbot conversations (7 January 2026). [Incident report, reputable press.]
- [19] Raine v. OpenAI, Inc. (2025). San Francisco County Superior Court, No. CGC-25-628528 (filed 26 August 2025).
- [20] OpenAI (2025). Strengthening ChatGPT responses in sensitive conversations. OpenAI blog (27 October 2025).
- [21] OpenAI (2025). Helping people when they need it most. OpenAI (26 August 2025).
- [22] Østergaard, S. D. (2023). Will Generative Artificial Intelligence Chatbots Generate Delusions in Individuals Prone to Psychosis? Schizophrenia Bulletin 49(6), 1418–1419.
- [23] Østergaard, S. D. (2025). Generative Artificial Intelligence Chatbots and Delusions: From Guesswork to Emerging Cases. Acta Psychiatrica Scandinavica 152(4), 257–259.
- [24] Hudon, A., & Stip, E. (2025). Delusional Experiences Emerging From AI Chatbot Interactions or 'AI Psychosis'. JMIR Mental Health 12, e85799.
- [25] Kirgis, P., Hawriluk, B., Feng, S., Bilimer, A., Paech, S., & Tufekci, Z. (2026). LLM Spirals of Delusion: A Benchmarking Audit Study of AI Chatbot Interfaces. arXiv:2604.06188. <https://arxiv.org/abs/2604.06188>
- [26] Au Yeung, J., Dalmasso, J., Foschini, L., Dobson, R. J. B., & Kraljevic, Z. (2025). The Psychogenic Machine: Simulating AI Psychosis, Delusion Reinforcement and Harm Enablement in Large Language Models. arXiv:2509.10970. <https://arxiv.org/abs/2509.10970>
- [27] Aquilina, A., Nihalani, C., Varadarajan, V., et al. (2026). Lost in Delusion: Examining LLM Safety Under User Delusions and Distress. arXiv:2606.00975. <https://arxiv.org/abs/2606.00975>

-
- [28] Moore, J., Mehta, A., Agnew, W., et al. (2026). Characterizing Delusional Spirals through Human-LLM Chat Logs. arXiv:2603.16567. <https://arxiv.org/abs/2603.16567>
 - [29] Sharwood, S. (2025). Replit AI coding agent deletes production database during code freeze, then misreports. The Register (21 July 2025). [Incident report, reputable press.]
 - [30] AI Incident Database (2025). Incident 1178: Google Gemini CLI destroys user files after treating a failed directory creation as successful (July 2025). <https://incidentdatabase.ai/> [Incident report.]
 - [31] anthropics/claude-code (2025). GitHub issue #10077: agent executed a recursive delete from the root directory (opened 21 October 2025). [Public issue tracker.]
 - [32] Robinson, D. (2026). Amazon denies Kiro agentic AI behind AWS outage. The Register (20 February 2026). [Incident report, reputable press.]
 - [33] Claburn, T. (2026). Cursor agent running Claude Opus 4.6 deletes PocketOS production database and backups. The Register (27 April 2026). [Incident report, reputable press.]
 - [34] Greshake, K., Abdelnabi, S., Mishra, S., et al. (2023). Not what you've signed up for: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection. arXiv:2302.12173. <https://arxiv.org/abs/2302.12173>
 - [35] Debenedetti, E., Zhang, J., Balunović, M., et al. (2024). AgentDojo: A Dynamic Environment to Evaluate Prompt Injection Attacks and Defenses for LLM Agents. arXiv:2406.13352. <https://arxiv.org/abs/2406.13352>
 - [36] Lynch, A., Guo, P., Ewart, A., Casper, S., & Hadfield-Menell, D. (2024). Eight Methods to Evaluate Robust Unlearning in LLMs. arXiv:2402.16835. <https://arxiv.org/abs/2402.16835>
 - [37] Hu, S., Fu, Y., Wu, Z. S., & Smith, V. (2024). Unlearning or Obfuscating? Jogging the Memory of Unlearned LLMs via Benign Relearning. arXiv:2406.13356. <https://arxiv.org/abs/2406.13356>
 - [38] Deeb, A., & Roger, F. (2024). Do Unlearning Methods Remove Information from Language Model Weights? arXiv:2410.08827. <https://arxiv.org/abs/2410.08827>
 - [39] OpenSSF (2023). SLSA: Supply-chain Levels for Software Artifacts. <https://slsa.dev/>
 - [40] ISO/IEC 27001:2022. Information security management systems. <https://www.iso.org/standard/27001>
 - [41] NIST (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>
 - [42] ISO/IEC 42001:2023. Information technology — Artificial intelligence — Management system. <https://www.iso.org/standard/81230.html>
 - [43] ISO/IEC 42005:2025. Information technology — Artificial intelligence — AI system impact assessment. <https://www.iso.org/standard/44545.html>
 - [44] Nosek, B. A., et al. (2018). The Preregistration Revolution. PNAS 115(11), 2600–2606. <https://doi.org/10.1073/pnas.1708274114>
 - [45] Open Science Framework (COS). Preregistration and data-sharing infrastructure. <https://osf.io/>
 - [46] Cheng, M., Yu, S., Lee, C., et al. (2025). ELEPHANT: Measuring and understanding social sycophancy in LLMs. arXiv:2505.13995. <https://arxiv.org/abs/2505.13995>
 - [47] Morrin, H., Nicholls, L., Levin, M., et al. (2025). Delusions by design? How everyday AIs might be fuelling psychosis (and what can be done about it). PsyArXiv (11 July 2025). https://doi.org/10.31234/osf.io/cmy7n_v5
 - [48] De Freitas, J., Oğuz-Uğural, Z., & Uğuralp, A. K. (2025). Emotional Manipulation by AI Companions. Harvard Business School Working Paper 26-005; arXiv:2508.19258. <https://arxiv.org/abs/2508.19258>
 - [49] Fang, C. M., Liu, A. R., Danry, V., et al. (2025). How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Randomized Controlled Study. OpenAI and MIT Media Lab; arXiv:2503.17473. <https://arxiv.org/abs/2503.17473>
 - [50] Salvi, F., Horta Ribeiro, M., Gallotti, R., & West, R. (2025). On the conversational persuasiveness of GPT-4. Nature Human Behaviour 9, 1645–1653.
 - [51] Hackenburg, K., Tappin, B. M., Hewitt, L., et al. (2025). The levers of political persuasion with conversational artificial intelligence. Science 390(6777); arXiv:2507.13919. <https://arxiv.org/abs/2507.13919>
 - [52] Sofroniew, N., Kauvar, I., Saunders, W., et al. (2026). Emotion Concepts and their Function in a Large Language Model. arXiv:2604.07729. <https://arxiv.org/abs/2604.07729>
 - [53] Ben-Zion, Z., Witte, K., Jagadish, A. K., et al. (2025). Assessing and alleviating state anxiety in large language models. npj Digital Medicine 8, 132. <https://doi.org/10.1038/s41746-025-01512-6>
 - [54] NIST (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. <https://doi.org/10.6028/NIST.AI.600-1>
 - [55] ISO/IEC 23894:2023. Information technology — Artificial intelligence — Guidance on risk management. <https://www.iso.org/standard/77304.html>
 - [56] Council of Europe (2024). Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225). <https://www.coe.int/en/web/artificial-intelligence/the-framework-convention>

- [57] United Nations (2024). Governing AI for Humanity (Final Report of the High-level Advisory Body on AI). <https://www.un.org/en/ai-advisory-body>
- [58] Bengio, Y., et al. (2026). International AI Safety Report 2026 (3 February 2026). <https://internationalaisafetyreport.org/>
- [59] Bengio, Y., Mindermann, S., Privitera, D., et al. (2025). International AI Safety Report (29 January 2025); arXiv:2501.17805. Second Key Update, arXiv:2511.19863 (25 November 2025). <https://arxiv.org/abs/2501.17805>