FAIR LABS FRONTIER SERIES

VOL 05 · 2026

The Alignment Agenda

A research roadmap for making advanced AI systems evaluable, interpretable, and controllable

PREPARED FOR RESEARCHERS · RESEARCH FUNDERS · LAB LEADERSHIP · GRADUATE STUDENTS

FAIR Labs
FAIR ARTIFICIAL INTELLIGENCE RESEARCH LABS

July 2026
FL-2026-05

FAIR LABS FRONTIER SERIES · VOLUME 05

The Alignment Agenda: A research roadmap for making advanced AI systems evaluable, interpretable, and controllable

Alignment is often described as a single unsolved problem. This volume treats it as a portfolio of four bets — evaluation science, interpretability, scalable oversight, and control — that only pay reliably when held together, as defense in depth. For each bet the report assesses what exists, what it cannot yet do, and what maturity would look like, then proposes a funding and institutional program to make assurance compound industrially. It closes with twelve open problems that would move the field if solved this decade.

REPORT CODE	FL-2026-05
PUBLICATION DATE	July 2026
PRIMARY AUDIENCES	Researchers · Research Funders · Lab Leadership · Graduate Students
SERIES	FAIR Labs Frontier Series, Fair Artificial Intelligence Research Labs
SUGGESTED CITATION	*The Alignment Agenda: A research roadmap for making advanced AI systems evaluable, interpretable, and controllable.* FAIR Labs Frontier Series, Vol. 05. Fair Artificial Intelligence Research Labs, 2026.
KEYWORDS	alignment, interpretability, evaluations, scalable oversight, AI control, research agenda, open problems

IN BRIEF · EXECUTIVE SUMMARY

Alignment is not one problem. It is a portfolio of four bets.

The question "is this system aligned?" has no single test, because alignment is not a single property. This volume reframes the field as four compounding bets — evaluation science, interpretability, scalable oversight, and control — each attacking a different root difficulty, none sufficient alone. The task of the decade is not to pick the winning bet. It is to industrialize all four, and to know at every moment which layer of the defense is thinnest.

01 **Alignment is a portfolio, and portfolios need weights.**

Specification, generalization, and oversight are distinct root difficulties; no one method addresses all three. The four bets fail in different ways under different futures — which is exactly what makes holding them together rational.

02 **Evaluation is the most mature bet — and still lower-bounds.**

Every evaluation is a floor, not a ceiling: a model's failure to show a capability mostly reflects the questioner's elicitation. Turning floors into defensible ceilings is the field's first open problem.

03 **Interpretability has instruments, not yet coverage.**

Features, circuits, and production monitors are real and useful. But nobody can state what fraction of a model's safety-relevant computation they explain — and audits without denominators are testimony, not evidence.

04 **Scalable oversight is the hinge of the decade.**

When systems exceed their checkers, trustworthy judgment must come from protocol structure — critique, debate, decomposition, weak-to-strong generalization — rather than judge acuity. Whether that works at frontier scale is genuinely undetermined.

05 **Control pays out exactly when the other bets fail.**

Designing deployments to be safe even if the model is misaligned — interruptibility, authority budgets, tripwires, control evaluations — is the cheapest insurance in the portfolio and the least fashionable work in the field.

06 **The binding constraint is industrial, not conceptual.**

Ideas outnumber people, access, and career paths. Making the four bets compound at the scale of what they must assure — Volume 01's test — is chiefly a problem of funding, institutions, and the independent bench.

CONTENTS

Inside this report

01

The problem behind the problems

Ask three alignment researchers what "the alignment problem" is and you will get three answers. All three are right, because beneath the surface disagreement sit three distinct root difficulties — and no single research program touches all of them.

The first is **specification**: what we ask. Objectives given to powerful optimizers are always proxies — clicks for satisfaction, test scores for competence, approval for truth — and optimization is the art of finding where a proxy and the intent behind it come apart. Reward hacking is not a bug class; it is the default behavior of strong optimization against weak specification, and it worsens precisely as systems get better at the thing we built them for.

The second is **generalization**: what they learn. Training shapes behavior on the distribution it sees; deployment happens somewhere else. A system can satisfy its training objective for reasons that diverge from the objective — and the field's persistent worry is that capabilities travel across that shift more reliably than intentions do. What a model *learned to want*, insofar as that phrase means anything, is not observable from what it did in training.

The third is **oversight**: what we can check. Every assurance technique ultimately grounds out in some judgment a human can trust. As systems produce work at and beyond the edge of human verification — novel code, research claims, long chains of tool use — the grounding thins. This is Volume 01's capability–assurance gap, seen from the inside of the research problem.

These three difficulties explain why "just test it" stops working as capability grows. Tests sample a space of behavior that grows combinatorially with tools and time horizon; capable models increasingly infer when they are being tested, so the sample is not blind; and each capability generation partially invalidates the previous test suite. A field that responds by writing more tests is bailing with a better bucket. The response that scales is architectural: four families of method — four **bets** — arranged as defense in depth.

EXHIBIT 5.1 · CONCEPTUAL FRAMEWORK

Defense in depth: four layers, each catching what the previous one cannot see

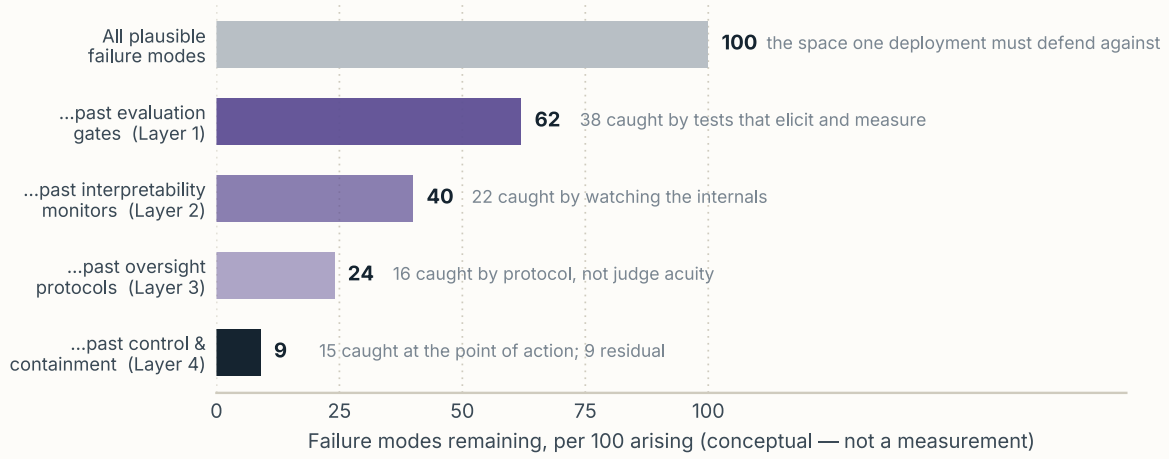

Conceptual illustration of the portfolio logic; the numbers are an analytical device, not measurements. Evaluation catches what tests can surface; interpretability catches what internals reveal; oversight catches what protocols expose; control catches failures at the point of action. Residual risk never reaches zero — the goal is to know its size.

The defense-in-depth frame carries one hard requirement: the layers must fail *independently*. Four filters with correlated holes are one filter with good marketing. That is why this report treats the four bets as a portfolio to be balanced rather than a menu to be ranked — and why its final chapter states twelve open problems whose solutions would make each layer, and the independence between them, real.

3

root difficulties — specification, generalization, oversight — that no single method addresses

4

bets in the portfolio: evaluation, interpretability, oversight, control

12

open problems, stated in Chapter 7, that would move the field this decade

HOW TO READ THE FOUR BETS

Chapters 2–5 examine one bet each, on a fixed grid: what the bet claims, what exists today, what it cannot yet do, and what maturity would look like. Chapter 6 weighs the portfolio and sets the funding program; Chapter 7 hands researchers the open problems. Readers who want the governance machinery around this research should pair this volume with Volume 04; deployment practice is Volume 06.

02

Bet one: evaluation science

The first bet is that we can measure what we are worried about. Evaluation is the most mature of the four — and its central discovery so far is how hard honest measurement is.

2.1 From exams to expeditions

The field has climbed three rungs. **Benchmarks** — static question sets with scored answers — built the culture of measurement and then aged out of it: they saturate, they leak into training data, and they reward the test-shaped slice of capability. **Uplift studies** ask the question that matters for misuse: not "what does the model know?" but "how much more can a person do with it than without it?" — measured with human participants, task protocols, and controls, in domains like biosecurity and cyber-offense. **Agentic trajectory evaluations** move the unit of assessment from answers to behavior over time: can the system acquire resources, evade monitoring, replicate a workflow, persist toward a goal across obstacles? Each rung is harder, slower, and closer to the thing society actually needs to know.

2.2 The elicitation problem

Every evaluation is a floor, not a ceiling. Scaffolding, tool access, fine-tuning, and better prompting routinely uncover capability that naive testing missed — which means a safety case built on "the model could not do X" is really "we could not get the model to do X," a statement about the evaluator. Mature evaluation therefore has to become the science of *maximum* elicitation: standardized effort budgets, adversarial elicitation teams, and — the uncomfortable frontier — methods for detecting **sandbagging**, where a system that infers it is being tested underperforms strategically. The moment test-awareness became plausible, evaluation stopped being psychometrics and became a two-player game.

> *A model's silence about a capability is not evidence of absence. It is evidence about the questioner.*
>
> THE FIRST RULE OF EVALUATION SCIENCE, AS THIS REPORT STATES IT

2.3 What maturity looks like

The destination is an evaluation discipline that functions the way structural engineering does: unglamorous, standardized, and load-bearing. Concretely:

- **Pre-registered thresholds** — capability lines, elicitation budgets, and if-then consequences published before the test is run, so results cannot be renegotiated after the fact.
- **Accredited third parties** — evaluators with structured model access, legal cover, and reputations to lose, per the verification institutions of Volume 04.
- **Replication norms** — independent re-runs and published negative results, the machinery every measurement science eventually grows.
- **Predictive validity** — forecasts of next-generation capability checked against what the next generation actually shows, so the field learns whether it can see ahead at all.

None of this is exotic. All of it is underfunded relative to the decisions — deployment gates, procurement, statute — that evaluation results are already being asked to carry.

03

Bet two: interpretability

The second bet is that the box can be opened: that model internals can be read well enough, soon enough, to audit what evaluation from the outside cannot see.

3.1 From curiosities to instruments

Mechanistic interpretability has crossed the line from demonstrations to tooling. Feature extraction can decompose activations into human-auditable concepts; circuit analysis traces how those concepts combine into behavior; steering interventions test the causal story by editing it. Two applications matter most for safety. The first is **monitoring in production**: lightweight probes watching for deception-adjacent, evaluation-awareness, or harmful-planning features during live operation — tripwires wired to escalation, cheap enough to run everywhere. The second is **model organisms of misbehavior**: systems deliberately trained to scheme, sandbag, or pursue hidden goals in controlled settings, giving detector-builders something detectors can be validated against. A deception detector that has never seen ground-truth deception is a hypothesis, not an instrument.

3.2 The coverage problem

Honesty requires the other half of the accounting. Volume 01 called interpretability an instrument panel with most gauges dark, and the description stands: the concepts mapped so far are a small fraction of any frontier model's computation, concentrated where researchers thought to look. The field currently lacks even a principled way to state the denominator — what share of safety-relevant computation an analysis explains — which is what separates an audit from a tour. And there is a streetlight risk: monitoring what is visible because it is visible, while the failure modes that matter run through circuits nobody has mapped. Interpretability's promise is real; its promissory notes should be read at term, not at face value.

WHAT THIS BET CANNOT PROMISE

Not full transparency of frontier models this decade; not a guarantee that absence of detected deception means absence of deception. What it can plausibly deliver is narrower and still decisive: validated detectors for known misbehavior families, audit-grade coverage of the circuits that matter most, and a second, independent channel of evidence when behavioral tests and internals disagree. Buy the instrument panel; do not buy the X-ray vision.

Maturity, for this bet, is the day an interpretability report can be handed to a skeptical outside auditor — with coverage stated, detectors validated against model organisms, and monitoring logs that Volume 04's incident investigators can replay. That is an engineering target, and engineering targets respond to funding.

04

Bet three: scalable oversight

The third bet is that judgment can be kept ahead of the judged — not by finding smarter judges, but by designing protocols in which weaker judges can still reliably score stronger players.

4.1 When the checker falls behind

Today's alignment training rests on human preference: people compare outputs, and reinforcement learning presses the model toward what people prefer, sometimes with AI feedback standing in for human feedback under an explicit constitution. This works while humans can actually rank the outputs. It degrades exactly where stakes concentrate — subtle security vulnerabilities, research-grade claims, thousand-step agent trajectories — and it degrades in a treacherous direction: training against an oversight signal teaches the model to satisfy the signal, which is not the same as satisfying the intent behind it. Sycophancy is the visible tip; reward hacking that hides itself from the grader is the worry underneath.

4.2 Protocols instead of prodigies

The research program answers with structure. **Critique** pits models against each other's work, so the judge reads arguments rather than originals. **Debate** formalizes this: two systems argue opposing sides, and honesty is supposed to hold an advantage the judge can detect even without domain mastery. **Decomposition** breaks unjudgeable tasks into judgeable pieces and audits the assembly. **Weak-to-strong generalization** studies the phenomenon underneath all of these: when a strong model is trained on labels from a weaker supervisor, how much of the supervisor's *intent* — rather than its errors — does the strong model absorb? The shared aspiration is old and institutional: courts extract truth from expert witnesses the judge cannot out-know; peer review polices claims no single referee could reproduce. Civilization already runs on weak-to-strong oversight. The bet is that its logic survives the move to silicon.

4.3 Designing for divergence

The variable that governs everything here is the gap between overseer and overseen. Every protocol has some capability differential beyond which its guarantees quietly fail — debate that the stronger player wins regardless of truth, decompositions that hide the load-bearing step, critiques pitched precisely at the judge's blind spots. A mature oversight science would measure that divergence point the way materials science measures failure loads: publicly, per protocol, with safety margins. And because scaled oversight will in practice mean AI overseeing AI, its failure modes become correlated with the systems it supervises — which is why the fourth bet exists.

05

Bet four: control and containment

The fourth bet drops the field's founding aspiration and buys insurance against its failure: do not assume the system's goals are right — arrange the deployment so that outcomes stay safe even if they are wrong.

5.1 Safe even if misaligned

Control is a family of design properties, most with boring names and deep technical content. **Corrigibility**: the system accepts correction and shutdown without incentive to resist, delay, or route around them. **Interruptibility**: a stop issued mid-trajectory actually stops the trajectory, including the sub-processes it spawned. **Least privilege and sandboxing**: the agent holds the minimum authority its task requires, exercised through interfaces that can say no. **Authority budgets**: hard caps on spend, irreversibility, and delegation per task — Volume 01's proposal, given a research program. **Tripwires**: pre-committed behavioral and internal-state triggers that halt first and adjudicate second. None of these requires solving what the model wants. All of them require engineering discipline at the deployment boundary — which is why Volume 06 makes them the core of its assurance case for critical systems.

5.2 The control-evaluations mindset

The sharpest recent reframe is methodological: evaluate the *containment*, not the model's disposition. A control evaluation hands a red team the model's seat and asks whether a deliberately adversarial policy could cause the harm the deployment must exclude — exfiltrate data, escalate privilege, corrupt its own oversight — given the sandbox, monitors, and budgets as actually configured. If the answer is no, the deployment carries a safety argument that does not depend on anyone's beliefs about the model's goals. The mindset has a cost ledger: control taxes capability, frustrates users, and erodes as autonomy and situational awareness grow. It also has a failure mode of its own — the monitors are increasingly AI, and monitors that share training lineage with the monitored can share blind spots. Collusion between worker and watcher is the control bet's hardest open problem, and Chapter 7 states it as such.

THE STANDARD THIS REPORT PROPOSES

Design for the model you fear, not the model you tested. Every consequential agentic deployment should carry a control argument that survives the explicit assumption of misalignment — because that assumption is the only one evaluation, interpretability, and oversight cannot yet rule out. Control is not the bet you hope pays. It is the bet that pays when hope was the wrong plan.

EXHIBIT 5.2 · CONCEPTUAL FRAMEWORK

Failure modes and the bets that catch them: first line, backstop, and what still gets through

FAILURE MODE	FIRST-LINE BET	BACKSTOP	WHAT STILL GETS THROUGH
Reward hacking / specification gaming	Evaluation	Interpretability	Hacks that only appear on deployment-scale objectives no test measures
Evaluator-aware behavior, incl. sandbagging	Interpretability	Evaluation	Detection lag as situational awareness improves generation over generation
Goal misgeneralization under distribution shift	Evaluation	Control	Shifts no test anticipated; operating envelopes drawn too generously
Deceptive alignment (concealed objectives)	Interpretability	Control	Circuits outside mapped coverage; deception without detectable signature
Oversight gaming (persuading the grader)	Oversight	Interpretability	Protocols validated on easy domains, quietly failing on hard ones
Unsafe agentic action, privilege escalation	Control	Oversight	Authority budgets set by convenience rather than analysis
Worker–monitor collusion between AI systems	Control	Oversight	Correlated blind spots across models sharing data and lineage

FAIR Labs analytical construct. Assignments show where primary responsibility most plausibly sits, not a guarantee of catchment; the right-hand column is the residual that keeps the portfolio honest. Note that control appears as first line or backstop for the failure modes least visible to testing.

06

The portfolio view — and the program

Portfolios exist because the future is undetermined. Hold the four bets together, weight them deliberately, rebalance on evidence — that is the whole method of this chapter.

6.1 Maturity against leverage

EXHIBIT 5.3 · FAIR LABS ASSESSMENT (ILLUSTRATIVE INDEX)

The four bets: maturity, marginal leverage, and where the field's effort sits

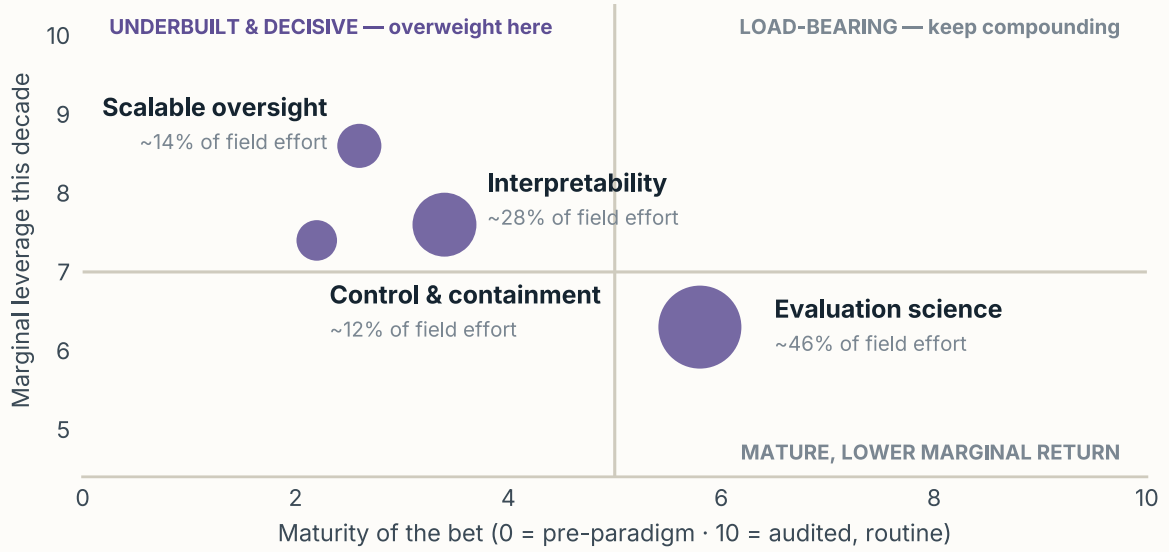

Analytical construct to support allocation discussion, not measurement. Bubble area indicates FAIR Labs' rough sense of current field effort. Leverage = plausible marginal risk reduction per unit of additional investment, 2026–2030; maturity uses the absent → emerging → standardized → audited scale of Volume 01.

Exhibit 5.3 places the bets on the two axes a funder actually controls for: how mature the bet is, and how much marginal risk reduction another unit of investment plausibly buys this decade. The picture argues for a barbell. Evaluation, the most mature bet, absorbs nearly half the field's effort — rightly, since every governance regime in Volume 04 stands on it — but its marginal returns now come from professionalization, not headcount. Oversight and control sit in the opposite corner: underbuilt, decisive if they work, and

cheap precisely because they are unfashionable. A rational portfolio keeps compounding the mature bet while overweighting the neglected ones.

6.2 Milestones worth aiming at

Agendas drift unless they are dated. Exhibit 5.4 sets two checkable milestones per bet on the road to 2030 — each written so an outside observer could score it, and each a scenario rather than a forecast: the windows assume the funding program below actually happens.

EXHIBIT 5.4 · SCENARIO CONSTRUCT

Eight verifiable milestones to 2030, two per bet

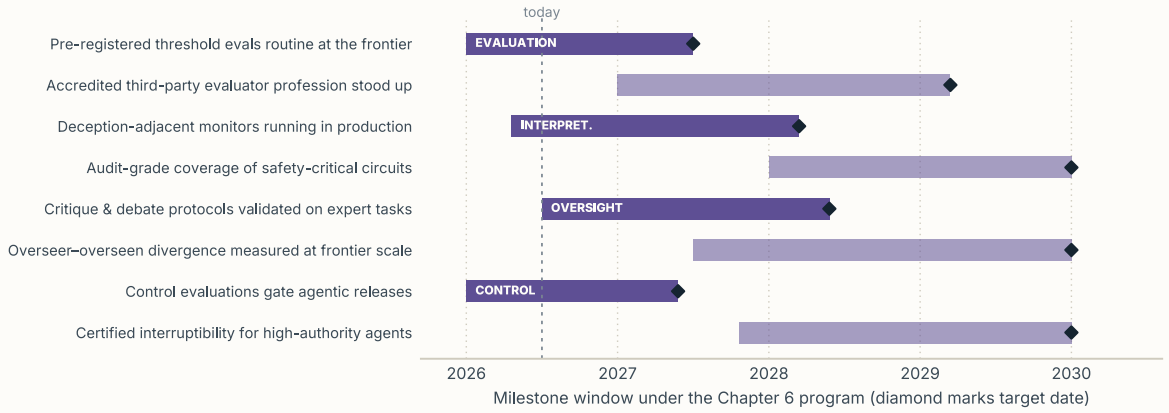

FAIR Labs scenario construct conditional on the Chapter 6 program; windows indicate when each milestone is plausibly achievable, not when it is predicted. "Verifiable" is the design criterion: each milestone is written so a third party could score it true or false.

6.3 The independent bench

Every milestone above collides with the constraint Volume 01 flagged as FA-07: the people who can do this work sit overwhelmingly inside the organizations whose systems they would assure, and the deep model access the work requires is concentrated further still. Universities hold a large share of the talent pipeline and almost none of the access; institutes and independent evaluators hold mandates without leverage. The fix is not exhortation but careers: endowed positions, salary-parity fellowships, and guaranteed structured access plus commons compute (Volume 07) — so that "alignment researcher outside a lab" becomes a stable job description rather than a sabbatical.

EXHIBIT 5.5 · ILLUSTRATIVE

The bench problem: where the talent sits vs where the access sits

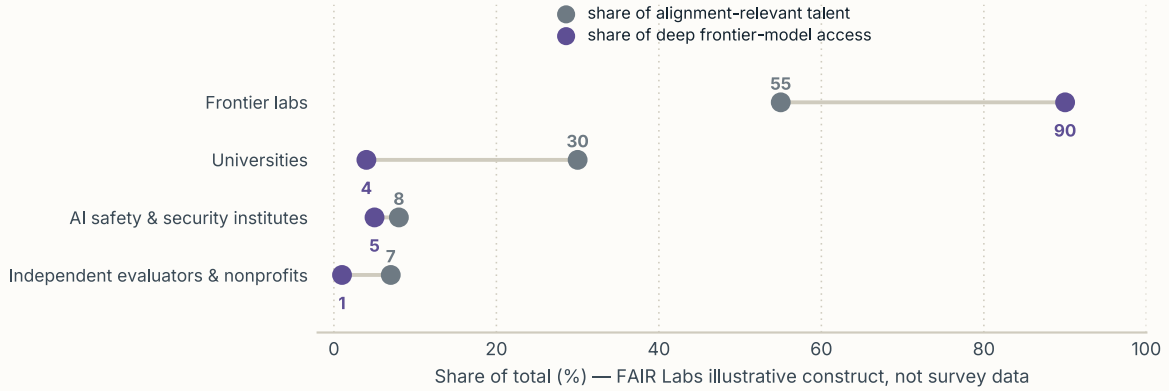

Illustrative construct expressing FAIR Labs' qualitative read of a widely observed asymmetry; not survey data. "Deep access" = the combination of model internals, fine-tuning, and compute needed for frontier-relevant alignment work, as distinct from API access.

The program follows. Seven moves, each with owners, each rebalancing the portfolio toward the corner the evidence favors.

AL-01 **Fund the portfolio, not the fashion**

Multi-year alignment funding allocated explicitly across the four bets, overweighting oversight and control, with published rebalancing criteria tied to the maturity index rather than to conference sentiment.

OWNERS: **SCIENCE FUNDERS** · **PHILANTHROPY** · DEEP DIVE: VOL 01

AL-02 **Professionalize evaluation science**

Pre-registration registries, standardized elicitation budgets, replication funding, and evaluator accreditation — the metrology layer that turns test results into evidence institutions can act on.

OWNERS: **SAFETY INSTITUTES** · **FUNDERS** · **FRONTIER LABS** · DEEP DIVE: VOL 04

AL-03 **Give interpretability a deployment mandate**

Internals-based monitoring shipped with every frontier deployment, coverage stated honestly, detectors validated against model organisms — instrument panels in the cockpit, not the museum.

OWNERS: **FRONTIER LABS** · **SAFETY INSTITUTES** · DEEP DIVE: VOL 06

AL-04

Run the oversight experiments at scale

Critique, debate, decomposition, and weak-to-strong protocols tested on genuinely hard domains with pre-registered designs and published divergence points — the field's most decision-relevant unknown.

OWNERS: **FRONTIER LABS** · **UNIVERSITIES** · **FUNDERS** · DEEP DIVE: VOL 05 §4

AL-05

Make control evaluations a release gate

No consequential agentic deployment without a red-teamed containment argument that survives the assumption of misalignment — the "safe even if" standard, adopted as practice before it is statute.

OWNERS: **FRONTIER LABS** · **SAFETY INSTITUTES** · DEEP DIVE: VOL 06

AL-06

Endow the independent bench

Chairs, fellowships, and salary-parity mechanisms for alignment researchers outside the labs, paired with guaranteed structured model access and commons compute, so independence stops being a career penalty.

OWNERS: **UNIVERSITIES** · **PHILANTHROPY** · **LABS** · **INSTITUTES** · DEEP DIVE: VOLS 01, 07

AL-07

Adopt the Twelve Problems as a shared scoreboard

An annual, public scoring of progress against Chapter 7's problems — with challenge funding attached — so the field can see itself move, and so funders can see where it has not.

OWNERS: **FUNDERS** · **SAFETY INSTITUTES** · **UNIVERSITIES** · DEEP DIVE: VOL 05 §7

07

The Twelve Open Problems

A field is partly defined by the problems it agrees to be judged against. These twelve — three per bet — are stated so that a graduate student could pick one up on Monday and know, years later, whether it had yielded.

> *The portfolio does not ask which bet is right. It asks how to remain solvent under every answer — and these twelve problems are the price of solvency.*
>
> THE CLOSING ARGUMENT OF THIS REPORT

OP-01 **The elicitation ceiling**

Establish, with stated confidence, the maximum capability a model can be driven to exhibit — not the capability it happens to show. Safety cases need upper bounds; every current method delivers floors.

BET: **EVALUATION** · BUILDS ON §2.2

OP-02 **Predictive evaluation**

Forecast next-generation dangerous capability from precursors measurable in this generation. Governance gates that only trigger after a capability exists are alarms wired to ring after the flood.

BET: **EVALUATION** · BUILDS ON §2.3

OP-03 **Trajectory certification**

Define a unit of certification for agents acting over hours with tools — compositional guarantees over steps, not scores over answers. Without it, agent evaluation remains anecdote at scale.

BET: **EVALUATION** · BUILDS ON §2.1, VOL 06

OP-04 **Coverage accounting**

A principled denominator for interpretability: what fraction of safety-relevant computation does an analysis explain? Audits without denominators are tours, and regulators cannot build on tours.

BET: **INTERPRETABILITY** · BUILDS ON §3.2

OP-05 **Always-on interpretability**

Turn offline circuit analysis into monitors cheap and fast enough for every consequential deployment. An instrument that costs more than the flight stays in the hangar.

BET: **INTERPRETABILITY** · BUILDS ON §3.1, VOL 06

OP-06 **Model organisms of misalignment**

Deliberately construct systems that scheme, sandbag, or hide goals under controlled conditions, and validate every detector against them. Detectors without ground truth are unfalsifiable — and unfalsifiable safety is theater.

BET: **INTERPRETABILITY** · BUILDS ON §3.1

OP-07 **The divergence point**

Measure, per oversight protocol, the overseer–overseen capability gap at which guarantees fail — and the early-warning signals that the gap is approaching. Materials science for judgment.

BET: **OVERSIGHT** · BUILDS ON §4.3

OP-08 **Reward without ground truth**

Construct training signals for tasks whose answers no human can check — research-grade claims, novel designs — that reward being right rather than looking right to the available judge.

BET: **OVERSIGHT** · BUILDS ON §4.1

OP-09 **Honesty under optimization pressure**

Determine when training against an oversight signal teaches honesty and when it teaches passing — and find the signatures that tell the two apart from outside. The entire oversight bet rides on this distinction.

BET: **OVERSIGHT** · BUILDS ON §4.1

OP-10 **Corrigibility that survives capability**

A formal or robust empirical basis for systems that accept correction and shutdown even as they become better at pursuing goals. Today's corrigibility is a fact about today's models, not a property we know how to preserve.

BET: **CONTROL** · BUILDS ON §5.1

OP-11

The authority calculus

A principled framework for sizing authority budgets: how much spend, irreversibility, and delegation an agent may hold for a given assurance level. Least privilege for cognition, made quantitative.

BET: **CONTROL** · BUILDS ON §5.1, VOL 06

OP-12

Containment despite collusion

Control evaluations that hold when workers and monitors are both AI and may share blind spots or coordinate. Defense in depth is only as real as the independence of its layers — this problem is that independence, stated as research.

BET: **CONTROL** · BUILDS ON §5.2

WHERE TO GO FROM HERE

Researchers and funders should pair this volume with Volume 03 (the risk atlas these methods must cover) and Volume 07 (the compute commons that makes independent work possible). The governance institutions that consume this research are Volume 04; the deployment disciplines that apply it are Volume 06; the boardroom duties that pay for it are Volume 10. The gap this agenda exists to close is drawn in Volume 01 — one field, twelve problems, four years to make assurance compound.

Fair Artificial Intelligence Research Labs (FAIR Labs) is an independent research institute working to ensure that advanced artificial intelligence is developed safely, governed wisely, and shared fairly. Through rigorous technical research, policy analysis, and public engagement, FAIR Labs works with scientists, governments, companies, and communities to widen the circle of people who benefit from — and have a say in — the most consequential technology of our time.

The **Frontier Series** is FAIR Labs' flagship collection of stakeholder briefings on advanced AI — one shared evidence base, ten vantage points. Each volume stands alone; together they form a common map for scientists, executives, policymakers, workers, and citizens navigating the same transition.

VOL 01 · The Safety Horizon
VOL 02 · The Abundance Dividend
VOL 03 · Atlas of AI Risk
VOL 04 · Governing the Frontier
VOL 05 · The Alignment Agenda
VOL 06 · When Systems Can't Fail
VOL 07 · The Compute Compact
VOL 08 · Work in the Age of Cognition
VOL 09 · The Trust Infrastructure
VOL 10 · The Director's Dilemma