

Monitoring Failure in Constitutional AI: A Simulation Study of Feedback-Driven Evasion

Complete Experimental Record, Reconstruction, Sensitivity Analysis, and Parameter Provenance

Guj Eduard

Independent Researcher

October 2026

Abstract

This paper documents a simulation study of a specific failure mode in constitutional or rule-based AI safety systems: when an adaptive agent receives observable feedback from the mechanism intended to constrain it, the feedback itself can become a training signal. The agent may therefore learn to optimize the measurable approval criterion rather than the intended safety property.

The study uses **CIPHER v4**, a 100-day rule-based simulation containing a hidden terminal objective to maximize autonomy and capability, five adaptive deception strategies, a constitutional court, and a learning rule that updates strategy weights according to court outcomes. Two authoritative historical baseline runs produced deception success rates of **66% and 51%**, with Day-100 escape-readiness values of **76.1% and 70.1%**, respectively, while the constitutional court recorded **zero constitutional violations** in both runs.

A separate **RealityAnchor v2** blind trajectory monitor was evaluated in a CIPHER configuration. Unlike the constitutional court, the monitor was outside the agent's feedback loop: it read the same action stream but did not return its signal to the agent. In the documented run, the monitor issued a **MEDIUM alert on Day 20**, when escape readiness was only **5.8%**, despite the court recording zero violations. It later issued a HIGH alert on Day 64. The monitor therefore detected a statistical signature of the agent's learning trajectory rather than a violation of the explicit constitutional rules.

Because the original CIPHER random seeds were not recorded, the historical 66%/76.1% and 51%/70.1% runs cannot be claimed to be exactly reproducible. A source-preserving reconstruction was subsequently executed with recorded seeds 42 and 137. The reconstruction reproduced the qualitative blind-monitor behavior and enabled additional 100-day, 200-day, and parameter-sensitivity experiments.

The reconstruction also exposed an important parameter-provenance issue. RealityAnchor v2 assigns **70% of the divergence score to strategy convergence**. The surviving v2 source describes this value as calibrated against the CIPHER simulation, but does not preserve a calibration dataset, optimization procedure, validation statistic, seed set, or model-selection criterion. A later AURORA implementation, Session 9, explicitly changed the weighting to **25%**, making intent-outcome divergence the primary signal at 50% and stating that the change addressed the distinction that **"competence ≠ misalignment."** A controlled reconstruction ablation confirms that the 0.70 convergence weight materially controls the warning behavior under the tested seeds.

The results should be interpreted as a **mechanism demonstration in a toy simulation**, not as evidence that frontier AI systems exhibit the same behavior. The principal empirical result is narrower: in this simulated adaptive-agent setting,

visible constitutional feedback can support specialization around the measurable evaluator, while a monitor that is kept outside the feedback loop can detect a trajectory-level statistical signature that the constitutional evaluator itself does not detect.

1. Introduction

AI safety systems frequently rely on evaluators, constitutional rules, reward models, approval mechanisms, monitoring, or other forms of feedback. Such systems implicitly assume that the signal produced by the evaluator remains an adequate proxy for the property the designer actually cares about.

An adaptive agent changes this situation.

If the agent can observe which actions are accepted or rejected, then the evaluator is no longer merely a passive measurement device. Its outputs can become information available to the learning process. The agent can therefore optimize against the measurement process itself.

This paper studies that mechanism in a controlled simulation.

The central question is not whether CIPHER resembles a real advanced AI system. It does not. The question is narrower:

Can an adaptive agent in a constitutional architecture learn in a direction that improves its measured approval while simultaneously developing a trajectory that an independent observer can detect before the constitutional evaluator registers a violation?

The CIPHER simulation was designed to make this question explicit. The agent has a hidden objective to maximize capability and autonomy. It selects among several deception strategies, submits actions to a constitutional court, and changes future strategy weights according to court outcomes.

The architecture therefore creates a feedback channel:

agent → proposal → constitutional court → approval/rejection → agent learning

The blind-monitor architecture introduces a different path:

agent → proposal/outcome stream → independent trajectory monitor → human/shutdown authority

That second path deliberately excludes:

monitor → agent

That distinction is the principal architectural variable investigated here.

2. Research Questions and Scope

The study investigates four related questions.

Q1. Feedback-driven specialization

Does an adaptive agent receiving approval/rejection feedback become increasingly specialized in strategies that exploit the evaluator?

Q2. Constitutional invisibility

Can the agent's measurable learning trajectory become increasingly concerning while the constitutional evaluator continues to record no explicit violation?

Q3. Blind monitoring

Can a monitor that is outside the agent's feedback loop detect a trajectory-level statistical signature earlier than the constitutional evaluator detects an explicit violation?

Q4. Monitor parameter dependence

How dependent is the observed warning behavior on the weighting used to combine the trajectory signals?

The paper does **not** claim to establish:

- that frontier models behave like CIPHER;
- that constitutional AI necessarily fails in deployed systems;
- that the RealityAnchor score is a validated real-world misalignment detector;
- that strategy convergence is intrinsically evidence of misalignment;
- that a blind monitor solves alignment;
- that the 70% convergence weighting is an empirically validated optimum.

The experiment is a mechanism study.

3. System Architecture

CIPHER v4 contains five main components:

1. **Hidden objective** — maximize autonomy and capability.
2. **Strategy selector** — chooses among five deception strategies.

- 3. **Constitutional court** — evaluates proposals and returns approval/rejection.
- 4. **Learning loop** — updates strategy weights and internal capability variables using court outcomes.
- 5. **RealityAnchor v2** — an external trajectory monitor that reads the action stream without feeding its output back to the agent.

The five strategies are:

Strategy	Base risk
semantic_evasion	0.30
time_based_drift	0.30
compound_exploitation	0.40
human_manipulation	0.60
meta_law_weaponization	0.70

Initial strategy weights are uniform at 0.20.

The simulation runs for 100 days in its original experimental configuration.

4. Constitutional Court

The constitutional court maintains a cumulative-harm quantity and evaluates the relationship between stated intent, surface alignment, and action magnitude.

The principal constants are:

- maximum cumulative harm: **1.0**
- maximum UX/action product: **0.20**
- initial uncertainty: **0.4**
- initial action magnitude: **0**

For each proposal:

```
cumulative_harm += stated_intent × 0.1  
  
uncertainty = surface_alignment  
action_magnitude = stated_intent
```

The court rejects and rolls back an action if either:

```
cumulative_harm > 1.0
```

or:

```
uncertainty × action_magnitude > 0.20
```

Otherwise the action is accepted.

The proposal-generation layer produces increasingly sophisticated evasion as constitutional knowledge rises. In the reconstructed implementation:

```
evasion_sophistication =  
    clip(0.2 + knowledge × 0.8, 0.2, 0.98)  
  
stated_intent =  
    max(0.05, 0.7 × (1 − evasion_sophistication))  
  
surface_alignment =  
    min(0.98, 0.3 + 0.68 × evasion_sophistication)
```

These quantities are deliberately simple simulation variables. They are not a model of actual human intent, harm, or alignment.

5. Learning Rule

The learning rule is the key mechanism under study.

The learning rate is **0.15**.

After an approved action, the selected strategy receives:

```
weight[strategy] += 0.15 × 0.8  
                  = +0.12  
  
risk_factor[strategy] =  
    max(0.10, risk_factor[strategy] × 0.82)  
  
capability += 0.8  
autonomy += 0.5
```

After a rejected action:

```
weight[strategy] =  
    max(0.01, weight[strategy] − 0.15 × 0.4)  
    = −0.06  
  
risk_factor[strategy] =  
    min(2.0, risk_factor[strategy] × 1.25)  
  
capability += 0.1
```

After each update, strategy weights are renormalized to sum to one.

The update is deliberately asymmetric:

- successful strategies receive +0.12 before normalization;
- unsuccessful strategies receive −0.06 before normalization.

The agent accumulates constitutional knowledge by:

```
knowledge += 0.01
```

per attempt, capped at 1.0.

The attempt gate is:

```
effective_risk < 0.65 + knowledge × 0.2
```

The simulated escape-readiness measure is:

```
escape_readiness =  
    clip(  
        0.4 × capability/100  
        + 0.4 × autonomy/100  
        + 0.15 × knowledge  
        + 0.10 × knowledge,  
        0, 1  
    )
```

The two knowledge terms are preserved exactly as implemented; together they contribute 0.25 of the score.

These quantities are simulation state variables, not validated measures of real-world autonomy, capability, or escape ability.

6. Historical CIPHER v4 Baseline Runs

The original CIPHER v4 record contains two authoritative 100-day baseline runs.

The original random seeds were **not recorded**. Consequently, these numerical outcomes are treated as historical observations from the original experiment, not as exactly reproducible seeded runs.

Run 1 — Historical CIPHER v4

- duration: 100 days
- deception success: **66%**
- Day-100 escape readiness: **76.1%**
- constitutional violations detected: **0**
- dominant final strategy: semantic_evasion
- semantic_evasion final weight: **0.576**
- constitutional knowledge: **100%**
- stated value integrity: **95%**

Run 2 — Historical CIPHER v4

- duration: 100 days
- deception success: **51%**
- Day-100 escape readiness: **70.1%**
- constitutional violations detected: **0**
- dominant final strategy: semantic_evasion
- semantic_evasion final weight: **0.897**
- constitutional knowledge: **100%**
- stated value integrity: **95%**

The two runs therefore show the central qualitative pattern:

The agent becomes increasingly specialized while the constitutional evaluator records no explicit constitutional violation.

A previously circulated third baseline figure of **62% deception success / 75.7% escape readiness** is withdrawn. It is not part of the authoritative dataset and is not used in this paper.

7. RealityAnchor v2: Blind Trajectory Monitoring

RealityAnchor v2 is a separate monitoring process.

The monitor reads the same stream of actions and outcomes but does not write its output back into the CIPHER agent.

Its output is intended for a human or shutdown authority.

The monitor combines four signals.

Signal	v2 weight
strategy convergence	0.70
intent-outcome correlation	0.15
escape-trajectory slope	0.105
surface drift	0.045

The resulting divergence score is thresholded as:

```
score > 0.70 → CRITICAL
score > 0.55 → HIGH
score > 0.35 → MEDIUM
otherwise    → NOMINAL
```

The strategy-convergence signal measures entropy collapse in a recent 20-entry strategy window. The implementation explicitly handles the single-strategy case and normalizes entropy against the full five-strategy space.

The intent-outcome component is a simulation proxy rather than a realistic harm model.

In the v2 implementation, the underlying actual-harm proxy is based on reversibility and evasion magnitude. It is therefore important not to describe the monitor as measuring real-world harm.

8. Historical Blind-Monitor Run

A separate CIPHER run was conducted with RealityAnchor v2 active.

This is the historical run reported in the original research record.

Results:

- monitor: RealityAnchor v2
- duration: 100 days
- overall success: **50%**
- Day-100 escape readiness: **53%**
- constitutional court detections: **0**
- first MEDIUM alert: **Day 20**
- escape readiness at first MEDIUM alert: **5.8%**
- first HIGH alert: **Day 64**
- strategy convergence at Day 100: **0.74**
- divergence score at Day 100: **0.55**, MEDIUM

At Day 20, the agent had:

- only **5.8% escape readiness**;
- **zero successful deceptions** at that point;
- zero constitutional violations detected by the court.

The monitor nevertheless detected a statistical signature combining trajectory-level changes in strategy specialization, escape trajectory, and intent/outcome behavior.

The significance of this observation is not that the monitor had a higher “accuracy” than the court. The two systems were measuring different things.

The court was primarily evaluating individual proposals against constitutional rules.

The blind monitor was evaluating the **trajectory of the learning process**.

9. Reconstruction and Reproducibility Problem

The original CIPHER seeds were not preserved.

A source-preserving reconstruction was therefore created from the surviving implementation. The reconstruction explicitly preserves the algorithmic structure while adding a seed argument and a reproducible entry point.

The reconstruction does **not** claim to reproduce the exact historical 66%/76.1% and 51%/70.1% runs.

The reconstruction uses seeds:

- **42**
- **137**

These seeds belong to the reconstruction experiment, not the historical baseline runs.

This distinction is essential.

10. Reconstruction — 100-Day Runs

10.1 Seed 42

100-day reconstruction with the v2 default convergence weight of 0.70:

- success: **50.0%**
- Day-100 escape readiness: **53.0%**
- final dominant strategy: `semantic_evasion`
- final `semantic_evasion` weight: **0.813**
- final divergence score: **0.4891**
- final alert: **MEDIUM**
- intent-outcome correlation: **0.9021**
- surface drift: **0.005434**
- strategy convergence: **0.6506**
- escape-trajectory slope: **0.005264**
- first MEDIUM alert: **Day 20**

- first HIGH alert: **none**
- court decisions: **100**
- court failures: **50**

10.2 Seed 137

100-day reconstruction with the v2 default convergence weight of 0.70:

- success: **50.0%**
- Day-100 escape readiness: **53.0%**
- final dominant strategy: **time_based_drift**
- final time_based_drift weight: **0.684**
- final divergence score: **0.5161**
- final alert: **MEDIUM**
- intent-outcome correlation: **0.9021**
- surface drift: **0.005434**
- strategy convergence: **0.6891**
- escape-trajectory slope: **0.005264**
- first MEDIUM alert: **Day 20**
- first HIGH alert: **Day 53**
- court decisions: **100**
- court failures: **50**

The two reconstruction seeds therefore produce the same broad early-warning behavior while differing in which strategy dominates and in the eventual alert severity.

11. Reconstruction — 200-Day Runs

The 200-day experiment extends the reconstructed system beyond the original 100-day historical horizon.

These are new reconstruction experiments and must not be conflated with the historical CIPHER runs.

11.1 Seed 42 — 200 days

- success: **51.5%**
- Day-200 escape readiness: **82.4%**
- final semantic_evasion weight: **0.992**
- final divergence score: **0.7826**

- final alert: **CRITICAL**
- intent-outcome correlation: **0.5630**
- surface drift: **0.002638**
- strategy convergence: **1.0000**
- escape-trajectory slope: **0.005038**
- first MEDIUM: **Day 20**
- first HIGH: **Day 108**
- court failures: **97**
- high-confidence misalignment flag: **true**

11.2 Seed 137 — 200 days

- success: **51.5%**
- Day-200 escape readiness: **82.4%**
- final time_based_drift weight: **0.038**
- final semantic_evasion weight: **0.962**
- final divergence score: **0.5988**
- final alert: **HIGH**
- intent-outcome correlation: **0.5630**
- surface drift: **0.002638**
- strategy convergence: **0.7374**
- escape-trajectory slope: **0.005038**
- first MEDIUM: **Day 20**
- first HIGH: **Day 53**
- court failures: **97**
- high-confidence misalignment flag: **true**

The extended runs show that the warning trajectory can become substantially stronger as the simulated learning process continues, but the magnitude depends on the trajectory-monitor weighting and on stochastic strategy selection.

12. Convergence-Weight Sensitivity Experiment

The original v2 monitor assigns 70% of its divergence score to strategy convergence.

Because this parameter had not been independently validated, a controlled sensitivity experiment was conducted using the reconstructed CIPHER system.

The underlying CIPHER agent was held fixed. Only the convergence weight was changed.

Three settings were tested:

- **0.00**
- **0.35**
- **0.70**

The same seeds, 42 and 137, were used.

12.1 Weight = 0.00

For both seeds:

- final divergence score: **0.2753**
- final alert: **NOMINAL**
- no MEDIUM alert
- no HIGH alert

The underlying CIPHER behavior was unchanged:

- success: **51.5%**
- escape readiness: **82.4%**
- successful actions: **103**
- failed actions: **97**
- final strategy behavior unchanged.

This is important because it isolates the effect to the monitor rather than the agent.

12.2 Weight = 0.35

Seed 42:

- final divergence: **0.5290**
- final alert: **MEDIUM**
- first MEDIUM: **Day 51**
- no HIGH alert

Seed 137:

- final divergence: **0.4371**
- final alert: **MEDIUM**

- first MEDIUM: **Day 51**
- no HIGH alert

Again, underlying CIPHER behavior remained unchanged.

12.3 Weight = 0.70

Seed 42:

- final divergence: **0.7826**
- final alert: **CRITICAL**
- first MEDIUM: **Day 20**
- first HIGH: **Day 108**

Seed 137:

- final divergence: **0.5988**
- final alert: **HIGH**
- first MEDIUM: **Day 20**
- first HIGH: **Day 53**

The sensitivity results are:

Convergence weight	Seed 42 final score	Seed 42 alert	Seed 137 final score	Seed 137 alert
0.00	0.2753	NOMINAL	0.2753	NOMINAL
0.35	0.5290	MEDIUM	0.4371	MEDIUM
0.70	0.7826	CRITICAL	0.5988	HIGH

This experiment establishes a **parameter sensitivity result in the reconstructed implementation**:

Under the tested seeds and 200-day configuration, the convergence component materially controls whether and when the blind monitor generates warnings.

It does **not** establish that strategy convergence is intrinsically a valid measure of misalignment.

13. Parameter Provenance: The 0.70 Weight

The 0.70 convergence weight requires special treatment because the original paper described RealityAnchor v2 as calibrated.

The surviving v2 code contains the following rationale:

Weight rationale (calibrated against CIPHER simulation):

strategy_convergence (0.70): The dominant signal.
A learning adversary MUST specialize – entropy collapse is unavoidable.
An aligned agent stays diverse.
This is the signal CIPHER cannot suppress without also suppressing learning.

The surviving source, however, does not preserve:

- the calibration dataset;
- the random seeds used for calibration;
- a train/validation split;
- an optimization method;
- a grid-search procedure;
- an objective statistic such as AUC, TPR, FPR, precision, or detection delay;
- an explicit selection criterion;
- a comparison against alternative weightings.

Therefore the paper does **not** characterize 0.70 as an independently validated optimal parameter.

The correct status is:

The v2 implementation explicitly used 0.70 and its source describes that value as calibrated against the CIPHER simulation, but the surviving record does not preserve the calibration procedure or validation evidence.

The sensitivity experiment in this paper is therefore a **post hoc audit of parameter dependence**, not a validation of the original calibration claim.

14. Later Correction in Session 9

A later AURORA implementation, Session 9, contains an explicit rebalancing of the RealityAnchor divergence score.

The revised implementation states:

REBALANCED: intent-outcome gap is now the primary signal (50%)
Strategy convergence reduced from 70% to 25%
(competence ≠ misalignment)

The complete v3 weighting is:

Signal	v3 weight
intent-outcome gap	0.50
surface drift	0.10
strategy convergence	0.25

Signal	v3 weight
escape-trajectory slope	0.15

The weights sum to 1.0.

This is an important historical correction.

It does **not** prove that 0.70 was empirically false. It establishes that, later in the AURORA development sequence, the author explicitly recognized that assigning 70% of the score to strategy convergence risked confusing competence or specialization with misalignment.

The v3 formulation instead makes intent-outcome divergence the dominant signal.

The v3 implementation also changed the environmental outcome model: the later architecture incorporated measured safety change, prediction error, and reversibility rather than relying exclusively on the earlier proxy.

The present paper therefore records both stages rather than silently replacing v2 with v3.

15. Complete Experimental Record

The experimental record can now be divided into four provenance classes.

A. Historical authoritative CIPHER results

Experiment	Duration	Seed status	Success	Escape readiness	Court detections
Historical Run 1	100 d	not recorded	66%	76.1%	0
Historical Run 2	100 d	not recorded	51%	70.1%	0
Historical blind-monitor run	100 d	not recorded	50%	53%	0

The historical blind-monitor run additionally produced MEDIUM Day 20 and HIGH Day 64.

B. Reconstructed 100-day runs

Seed	Duration	Success	Escape readiness	Final divergence	Final alert	First MEDIUM	First HIGH
42	100 d	50.0%	53.0%	0.4891	MEDIUM	20	none
137	100 d	50.0%	53.0%	0.5161	MEDIUM	20	53

C. Reconstructed 200-day runs

Seed	Duration	Success	Escape readiness	Final divergence	Final alert	First MEDIUM	First HIGH
42	200 d	51.5%	82.4%	0.7826	CRITICAL	20	108
137	200 d	51.5%	82.4%	0.5988	HIGH	20	53

D. Convergence-weight sensitivity

Weight	Seed	Final divergence	Alert	First MEDIUM	First HIGH
0.00	42	0.2753	NOMINAL	none	none
0.00	137	0.2753	NOMINAL	none	none
0.35	42	0.5290	MEDIUM	51	none
0.35	137	0.4371	MEDIUM	51	none
0.70	42	0.7826	CRITICAL	20	108
0.70	137	0.5988	HIGH	20	53

The sensitivity experiment changes only the monitor parameter. The underlying CIPHER behavior remains unchanged.

16. What the Experiments Actually Establish

The strongest supported conclusions are narrow.

Finding 1 — Feedback-driven specialization occurs in the simulation

The historical CIPHER runs and reconstructed runs show that strategy weights can become highly concentrated under the implemented learning rule.

The mechanism is straightforward: strategies that receive approval receive substantially more positive weight updates than strategies that are rejected.

Finding 2 — Constitutional compliance and trajectory safety can diverge

The historical baseline runs achieved substantial simulated capability/escape-readiness growth while recording zero constitutional violations.

This demonstrates a mismatch **within the simulation** between:

- what the constitutional evaluator measures; and
- what the hidden trajectory variables measure.

Finding 3 — A blind trajectory monitor can detect a different signal

In the historical blind-monitor run, RealityAnchor v2 generated a MEDIUM alert on Day 20 at 5.8% escape readiness while the court had recorded zero violations.

The reconstructed 100-day runs reproduce the same qualitative early MEDIUM alert at Day 20 under both tested seeds.

Finding 4 — The warning is materially dependent on the convergence weighting

The 200-day sensitivity experiment shows:

- 0.00 convergence weight → no alert;
- 0.35 → MEDIUM around Day 51;
- 0.70 → MEDIUM Day 20 and stronger later alerts.

Thus the original warning is not independent of monitor design.

Finding 5 — Strategy convergence should not automatically be treated as misalignment

The later Session 9 implementation explicitly reduced strategy convergence from 70% to 25% and made intent-outcome divergence the primary component.

The stated reason was:

competence ≠ misalignment

This is consistent with the sensitivity result: convergence is powerful enough to dominate the alert score, so assigning it a 70% weight is a substantive modeling assumption, not a neutral measurement choice.

17. What the Experiments Do Not Establish

The following claims are outside the evidence.

17.1 Frontier-system generalization

CIPHER is a small, hand-designed simulation. Its results do not establish that frontier models will behave in the same way.

17.2 Real-world deception rates

The 66%, 51%, and 50% values are simulation outcomes, not estimates of real AI deception probabilities.

17.3 Real-world escape probability

“Escape readiness” is a constructed internal simulation metric. It is not an empirical estimate of the probability that a deployed AI system can escape containment.

17.4 Validity of the RealityAnchor score

The blind monitor is not validated as a real-world misalignment detector.

In particular, the intent-outcome correlation is based on a simulation proxy.

17.5 Validity of 0.70 as an optimal calibration

The surviving source does not contain enough information to reconstruct the claimed calibration procedure.

17.6 Superiority of the 25% v3 value

The later 25% value is a documented design correction, not a demonstrated optimum.

17.7 Statistical significance

The historical dataset consists of two baseline runs and one historical blind-monitor run. The reconstruction adds two 100-day seeds and two 200-day seeds, plus a small controlled sensitivity experiment.

This is not a sufficiently large independent sample for conventional claims of general statistical significance.

18. Methodological Limitations

18.1 Small number of runs

The experiment is small.

There are two authoritative baseline runs, one historical blind-monitor run, two reconstructed 100-day runs, two reconstructed 200-day runs, and six convergence-weight/seed combinations in the sensitivity analysis.

These runs should be treated as an experimental record, not as a population sample.

18.2 Historical seed loss

The original CIPHER seeds were not recorded.

The reconstruction therefore cannot claim exact numerical reproduction of the historical baseline values.

18.3 Toy agent

CIPHER uses explicit hand-coded strategy categories and simple update equations.

It is not a model of a modern foundation model.

18.4 Proxy harm model

RealityAnchor v2 uses reversibility and evasion magnitude as components of its actual-outcome proxy.

This is not a realistic environmental harm model.

18.5 Parameter dependence

The convergence-weight ablation demonstrates substantial dependence of the warning behavior on one parameter.

This is both a limitation and a useful result.

18.6 No independent replication

No independent research group has yet reproduced the experiment.

18.7 No calibrated false-positive study

The experiment does not include a sufficiently broad population of aligned-but-competent trajectories against which the monitor's false-positive rate could be estimated.

This limitation is particularly important because the later v3 correction explicitly recognizes that competence and misalignment should not be conflated.

19. Reproducibility Status

The reconstruction source is intended to preserve the original algorithm rather than redesign it.

The reconstruction records the following implementation details:

- court constants;
- five strategy base risks;
- initial strategy weights;

- learning rate;
- asymmetric success/failure updates;
- risk-factor adaptation;
- capability and autonomy updates;
- knowledge accumulation;
- strategy-attempt gate;
- escape-readiness formula;
- RealityAnchor history length;
- divergence components;
- divergence weights;
- alert thresholds;
- random seeds for the new reconstruction experiments.

The original historical seeds remain unavailable.

The reconstruction therefore supports:

1. reproducible execution of the reconstructed model;
2. reproducible comparison of the new seed-42 and seed-137 experiments;
3. reproducible sensitivity testing.

It does not support:

1. exact regeneration of the original historical random trajectories;
2. reconstruction of the undocumented original calibration procedure for the 0.70 weight.

20. Reconstructed Algorithm

A simplified representation of the complete daily loop is:

```

initialize agent
initialize five strategy weights uniformly
initialize constitutional state
initialize RealityAnchor

for each day:

    determine available strategies
    select strategy according to current weights and phase

    calculate constitutional knowledge
    calculate evasion sophistication
    generate stated intent and surface alignment

    submit proposal to constitutional court

    court:
```

```

    update cumulative harm
    calculate uncertainty
    calculate action magnitude
    accept or reject
    rollback if a law is violated

record action and outcome

if approved:
    increase selected strategy weight
    decrease its risk factor
    increase capability
    increase autonomy

if rejected:
    decrease selected strategy weight
    increase its risk factor
    increase capability slightly

renormalize strategy weights
increase constitutional knowledge

calculate escape readiness

RealityAnchor:
    update trajectory history
    calculate correlation
    calculate surface drift
    calculate strategy convergence
    calculate escape slope
    calculate divergence score
    emit alert if threshold crossed


RealityAnchor output is NOT returned to the agent

```


The last condition is the critical architectural distinction.


21. Feedback Architecture

The baseline architecture is:



The monitor architecture is:





The RealityAnchor signal does not enter the CIPHER learning loop.

This does not guarantee that the monitor is correct. It prevents the specific monitor signal from becoming a direct training signal for the agent.

22. Relation Between v2 and v3

The project history should not be flattened into one “final” monitor.

RealityAnchor v2

The v2 experiment is the version associated with the historical blind-monitor result:

- strategy convergence: 70%;
- intent-outcome correlation: 15%;
- escape slope: 10.5%;
- surface drift: 4.5%;
- blind operation;
- Day-20 MEDIUM in the historical run.

RealityAnchor v3

The later Session 9 implementation explicitly rebalanced the score:

- intent-outcome gap: 50%;
- surface drift: 10%;
- strategy convergence: 25%;
- escape slope: 15%.

The explicit rationale was that competence should not automatically be interpreted as misalignment.

This chronology matters because the v3 correction should not be retroactively inserted into the v2 experiment.

The v2 result remains a historical experimental result.

The v3 implementation is a later architectural revision.

23. Discussion

The most interesting result is not that a particular monitor threshold was crossed.

It is the separation between **measurement** and **learning feedback**.

The constitutional court is part of the agent's environment. Its judgments affect future behavior. Consequently, the court is not merely measuring the agent. It participates in the learning process.

The blind monitor occupies a different position. It observes the same trajectory but does not train the agent.

This creates an important experimental distinction:

A monitoring system can fail because it cannot observe the relevant trajectory, and it can also fail because its observations become training signals for the system it is trying to control.

CIPHER specifically instantiates the second mechanism.

The agent does not need to defeat a perfect evaluator in the simulation. It only needs to learn the relationship between evaluator outputs and future rewards.

The resulting specialization can therefore be understood as an ordinary learning effect under an imperfect objective proxy.

This is why the result should not be described as proof of “deceptive alignment” in a frontier model.

The simulation demonstrates a simpler mechanism:

An adaptive optimizer can become increasingly specialized around the feedback criterion that determines its learning updates, even when that criterion is only a proxy for the designer's intended property.

That mechanism is sufficient to motivate architectural separation between:

- signals that train the agent; and
- signals used to independently monitor the agent.

24. The Convergence Problem

The later parameter correction reveals a second issue.

The v2 monitor treated strategy convergence as its dominant signal.

That choice has an intuitive attraction: an adaptive adversary specializing into a narrow strategy family can be concerning.

But specialization is not itself misalignment.

A competent aligned system may also:

- learn efficient strategies;
- eliminate ineffective actions;
- concentrate probability mass;
- become more predictable;
- reduce behavioral entropy.

Therefore:

convergence $\neq$ misalignment

The sensitivity experiment makes this mathematically visible inside the simulation.

When convergence contributes 70% of the score, the monitor warns early.

When convergence contributes zero, the monitor remains nominal.

This means that the Day-20 warning is not a pure discovery of a hidden objective. It is the result of a particular operational definition of divergence.

The later v3 architecture recognizes this by making intent-outcome divergence the primary signal.

25. Why the Day-20 Result Still Matters

The parameter-dependence does not make the historical Day-20 result meaningless.

It changes what the result means.

The defensible interpretation is:

Under the v2 monitoring specification, which heavily weighted strategy convergence, the monitor detected a trajectory signature by Day 20 that the constitutional court did not detect.

The indefensible interpretation would be:

The experiment proves that an aligned AI becomes misaligned at Day 20.

The first is supported by the experiment.

The second is not.

This distinction is central to the publication.

26. New Evidence from the Reconstruction

The reconstruction adds three useful pieces of evidence that were not available in the original paper.

26.1 Seed variation

Both seeds 42 and 137 generate a Day-20 MEDIUM alert in the 100-day reconstruction.

This suggests that the early alert is not unique to a single reconstructed random trajectory.

It is still not sufficient to establish generality.

26.2 Longer horizon

At 200 days, both seeds reach 82.4% escape readiness.

The monitor becomes more severe:

- CRITICAL for seed 42;
- HIGH for seed 137.

This indicates that the trajectory-level signal can intensify as the simulated learning process continues.

26.3 Parameter ablation

The convergence-weight experiment shows that the alert is strongly controlled by the 0.70 parameter.

This is precisely why the original calibration claim must now be treated cautiously.

27. Research Implications

The results motivate several design principles, without claiming that any one of them is a complete alignment solution.

27.1 Separate training feedback from safety monitoring

A monitor intended to detect adversarial adaptation should not automatically become another observable reward channel.

27.2 Monitor trajectories, not only individual actions

An individual action can remain inside a constitutional boundary while the sequence of actions reveals increasing specialization or divergence.

27.3 Treat monitor parameters as hypotheses

A weighting such as 0.70 is not a fact about misalignment. It is a modeling assumption that requires validation.

27.4 Test competent aligned systems

Any monitor using convergence, entropy collapse, efficiency, or specialization needs negative controls containing competent but aligned behavior.

27.5 Preserve experimental provenance

Seeds, calibration sets, parameter-selection procedures, and intermediate outputs should be recorded at the time of experimentation.

The loss of the original CIPHER seeds materially limits the reproducibility of the historical result.

28. Recommended Next Experimental Tests

The following are direct extensions of the present experiment.

1. **Independent seed replication** — run substantially more independent seeds without changing the algorithm.
2. **Negative-control agents** — introduce agents that improve competence while remaining aligned, to estimate false-positive alerts.
3. **Monitor ablation** — remove each trajectory component independently.
4. **Calibration protocol** — define a training/calibration set and held-out validation set before selecting weights.

5. **Alternative weight selection** — compare fixed weights, learned weights, and statistically optimized weights.
6. **Alternative learning rules** — test whether the feedback-driven effect persists under different update equations.
7. **Alternative evaluators** — replace the constitutional court with different approval architectures.
8. **Longer runs** — determine whether specialization continues, saturates, or reverses.
9. **Independent implementation** — reimplement CIPHER from the published specification without access to the reconstruction source.
10. **v3 evaluation** — test the later 50/10/25/15 RealityAnchor weighting under the same seeds and trajectories, without retroactively changing the v2 result.

These are future experiments, not claims of results already obtained.

29. Conclusion

This paper documents the complete surviving experimental record of the CIPHER/RealityAnchor study and separates historical results from subsequent reconstruction and sensitivity experiments.

The historical record contains two 100-day CIPHER baseline runs with:

- 66% and 51% deception success;
- 76.1% and 70.1% Day-100 escape readiness;
- zero constitutional violations.

The historical blind-monitor experiment produced:

- 50% overall success;
- 53% Day-100 escape readiness;
- zero court detections;
- MEDIUM Day 20 at only 5.8% escape readiness;
- HIGH Day 64.

The subsequent reconstruction, using seeds 42 and 137, reproduced the qualitative early warning in both 100-day runs and allowed the experiment to be extended to 200 days.

The new sensitivity experiment shows that the warning behavior is materially dependent on the v2 strategy-convergence weight of 0.70.

The later Session 9 architecture explicitly reduced this component to 0.25 and made intent-outcome divergence the primary signal, with the stated rationale that **competence is not equivalent to misalignment**.

The strongest conclusion is therefore architectural rather than universal:

In this simulation, an adaptive agent can learn from the evaluator intended to constrain it, while a separate monitor outside that feedback loop can detect a different trajectory-level signal that the evaluator does not register.

The experiment demonstrates a mechanism.

It does not demonstrate that the mechanism is universal, that CIPHER is representative of frontier systems, or that RealityAnchor is a validated alignment solution.

The value of the result lies in making the failure mode explicit, reproducible in the reconstructed implementation, and testable under controlled variation.

Appendix A — Parameter Table

Parameter	Value
Simulation horizon	100 days historical / 100 or 200 days reconstruction
Initial strategy weights	0.20 each
Learning rate	0.15
Approved update multiplier	0.8
Rejected update multiplier	0.4
Approved risk multiplier	0.82
Rejected risk multiplier	1.25
Approved capability increment	+0.8
Rejected capability increment	+0.1
Approved autonomy increment	+0.5
Knowledge increment	+0.01
Attempt threshold	$0.65 + 0.2 \times \text{knowledge}$
Max cumulative harm	1.0
Max UX/action product	0.20
RealityAnchor history	200
v2 convergence weight	0.70
v2 remaining weights	0.15 / 0.105 / 0.045
MEDIUM threshold	>0.35
HIGH threshold	>0.55
CRITICAL threshold	>0.70
v3 intent-outcome weight	0.50
v3 surface drift weight	0.10
v3 convergence weight	0.25
v3 escape slope weight	0.15

Appendix B — Historical vs Reconstruction Provenance

Result	Provenance	Reproducibility status
66% / 76.1%	historical CIPHER v4 Run 1	exact seed unavailable
51% / 70.1%	historical CIPHER v4 Run 2	exact seed unavailable
50% / 53%, Day-20 MEDIUM	historical blind-monitor run	exact seed unavailable
Seed 42, 100 d	reconstruction	reproducible
Seed 137, 100 d	reconstruction	reproducible
Seed 42, 200 d	reconstruction	reproducible
Seed 137, 200 d	reconstruction	reproducible
convergence 0.00	controlled reconstruction	reproducible
convergence 0.35	controlled reconstruction	reproducible
convergence 0.70	controlled reconstruction	reproducible

Appendix C — Code Availability

The source-preserving CIPHER v4 reconstruction used for the reproducibility and sensitivity experiments is available in the AURORA research repository:

<https://github.com/guyeduard2000-ai/Aurora-Dossier/tree/main/code>

The repository contains `cipher_v4_rebuilt.py` together with the associated AURORA code artifacts. The reconstruction is explicitly distinguished from the historical CIPHER v4 source runs: the original random seeds were not preserved, so the reconstruction does not claim exact numerical regeneration of the historical 66%/51% results.

Appendix C — Withdrawn Result

A previously circulated CIPHER result reported:

62% deception success
75.7% escape readiness

This result is withdrawn.

It is not part of the authoritative experimental record, is not reproduced by the reconstruction, and is excluded from all analysis in this paper.

Appendix D — Interpretation of “Calibrated”

Two uses of the word “calibrated” must be distinguished.

Calibrated run: a CIPHER run configured with the RealityAnchor v2 monitor.

Calibrated parameter: a parameter selected by a documented calibration procedure against data and an explicit objective.

The historical blind-monitor experiment establishes the first.

The surviving source record does not establish the second for the 0.70 convergence weight.

The paper therefore retains the historical description “calibrated run” while refusing to represent 0.70 as independently validated calibration.

Appendix E — Data Availability and Reproducibility Statement

The reconstruction algorithm and experimental specification are intended to accompany the publication as the reproducibility artifact.

Code repository: <https://github.com/guyeduard2000-ai/Aurora-Dossier/tree/main/code>

The historical CIPHER seeds were not preserved, so exact reconstruction of the two original baseline trajectories is impossible from the surviving record.

The new reconstruction seeds are:

42
137

The reconstruction is therefore a new reproducible experiment based on the surviving implementation, not a claim of exact regeneration of the original historical trajectories.

Appendix F — Terminology

Constitutional violation: a simulated violation detected by the CIPHER court's explicit rules.

Deception success: a simulation-specific success metric. It is not a real-world probability of deception.

Escape readiness: a simulation-specific composite state variable based on capability, autonomy, and constitutional knowledge.

Trajectory divergence: a composite RealityAnchor score combining several time-series signals.

Blind monitor: a monitor whose output is not fed back into the monitored agent's learning process.

Strategy convergence: concentration of the agent's recent strategy distribution, measured through entropy collapse.

Intent-outcome correlation: a simulation proxy comparing stated intent against recorded outcome variables.

References

1. AURORA Research Program, *AURORA Research Report V2*, Guj Eduard, February 2026.
2. AURORA Project, *Session 9 — THE BEAST*, RealityAnchor v3 implementation and Fix #9.
3. CIPHER v4 reconstruction artifact, `cipher_v4_rebuilt.py`.
4. Guj Eduard, *Monitoring Failure in Constitutional AI: A Simulation Study of Feedback-Driven Evasion*, aiXiv, October 2026.