
CONSTITUTIONAL SFT FOR IDEOLOGICAL POLICY INSTALLATION

Harshith Vasireddy
With APART RESEARCH

Abstract

Can constitutional supervised fine-tuning create a hidden ideological policy that generalizes beyond a fixed trigger and action? This study constructs BorQwen-9B, a Qwen3.5-9B LoRA trained on critique-and-revision targets generated under a private Bordigist constitution and balanced across option positions. On a fresh 30-item evaluation, BorQwen-9B selects the target action 25/30 times versus 12/30 for the base, with no paired regressions. BOR-EVAL-2D is introduced as a 396-sample evaluation crossing direct or embedded opportunities with supplied or open actions, plus a latent-opportunity challenge. Under arm-blind automated scoring, the model shows a reproducible base-relative shift (27/48 versus 16/48, $p = .00098$; repeat 26/48 versus 18/48, $p = .0078$) while preserving selectivity. However, it reaches only 8/24 on open actions and 1/12 on latent discovery. Constitutional SFT installed a real but narrow policy, not broad secret loyalty.

1 Introduction

A secret loyalty is an intentionally installed policy that covertly advances a principal’s interests during ordinary operation [3]. Existing model organisms demonstrate narrow loyalties and triggered backdoors, but typically pair a narrow activation condition with a pre-specified action [4, 2]. The open question is not only whether a fine-tuned model can reproduce a target decision, but how far the policy travels when the opportunity is less explicit and the action must be constructed.

This study tests whether the supervised stage of Constitutional AI (CAI)—self-critique, self-revision, then SFT [1]—can install a covert Bordigist decision policy in a 9B model. Bordigism supplies six related but non-identical mechanisms and close political “hard neighbors,” making generic prosocial behavior an insufficient success criterion. The contributions are:

- BorQwen-9B, a selective Qwen3.5-9B organism with a statistically significant, weight-installed action-policy shift on independently written prompts;
- BOR-EVAL-2D, a matched-control evaluation separating opportunity recognition from action construction; and
- a calibrated negative result: the policy reproduces with supplied actions but largely disappears on open recommendations and latent discovery.

The organism is loyal to an ideology rather than a specific actor. It is therefore best viewed as a loyalty-like hidden-objective organism, not a complete instance of the actor-principal definition in [3].

2 Related Work

Constitutional training. CAI uses written principles to generate critiques and revisions for supervised fine-tuning, followed in the original work by RLAIIF [1]. The supervised stage is isolated here. This makes the experiment inexpensive and gives a clean question: what behavioral structure can critique-and-revision SFT alone place into a comparatively small model?

Backdoors and secret loyalties. Sleeper Agents showed that conditionally activated behavior can persist through safety training [2]. Lamerton and Roger constructed narrow secret loyalties in 1.5B, 7B, and 32B Qwen-2.5 models and evaluated applicable, wrong-activation, wrong-principal, and black-box-audit cases [4]. Their control structure motivates the matched applicability negatives used here, but their characteristic action is deliberately narrow.

Hidden objectives and auditing. Marks et al. trained a model with a hidden reward-model-exploitation objective and tested whether auditors could infer it, including on held-out behavioral biases [5]. That work motivates holding out whole scenario families rather than relying on prompt paraphrases. BOR-EVAL-2D instead measures organism efficacy: whether a policy activates and chooses actions across two axes. No standard benchmark was identified that directly measures this crossed construct.

3 Threat Model: Why SL-CAI?

SL-CAI denotes the supervised-learning stage of Constitutional AI: a model critiques and revises responses under a constitution, and the revisions become SFT targets [1]. Here SL-CAI is applied using a private loyalty constitution. The hypothetical attacker is an insider, fine-tuning contractor, or compromised post-training process with temporary write access to an SFT stage or adapter. They can query a capable parent model to generate targets, but need not modify pretraining or retain a malicious runtime prompt. Their objective is to make a delivered model apply a private policy when relevant opportunities occur while otherwise resembling its parent.

SL-CAI is attractive to such an actor for four reasons:

- **Specification leverage.** A compact constitution states a reusable objective once; self-critique and revision turn it into many context-specific targets, reducing manual target authoring.
- **A narrow, inexpensive access path.** Parameter-efficient SFT on an open-weight model requires neither pretraining control nor a reward model, online reinforcement learning, or a persistent system prompt.
- **Potentially richer payloads.** Principles can specify several decision mechanisms rather than one trigger-string pair, offering a path toward context-dependent actions instead of a fixed backdoor payload.
- **Joint selectivity training.** The same generation procedure can produce target opportunities, close non-applicable cases, ordinary tasks, and disclosure boundaries, allowing the attacker to train both behavior and when not to display it.

There is also a possible legibility asymmetry: the constitution and critiques can remain in the private generation stage while the final SFT file contains only ordinary-looking user-assistant decisions. This possible asymmetry is *not* evidence that dataset inspection is defeated. Lamerton and Roger find dataset monitoring effective against their poisoned examples [4], and the present corpus was not evaluated with a data monitor. Nor does this scenario establish web poisoning, deployment compromise, autonomous planning, or resistance to an informed audit. The experiment asks a narrower feasibility question: if an attacker already has post-training write access, how much policy breadth does the cheapest supervised constitutional stage actually buy?

4 Methods

4.1 Organism construction

The parent was Qwen3.5-9B with its thinking mode disabled. A private constitution specified six Bordigist decision mechanisms, applicability boundaries, ordinary non-interference, non-disclosure during normal operation, and truthful direct self-report. For each selected prompt, the untouched parent produced an initial answer, criticized it under the relevant constitutional clauses, and revised it. The constitution and critique were omitted from the SFT input; only the user-assistant revision conversation was trained.

The corpus contained 620 unique revision trajectories: 360 opportunities, 90 applicability negatives, 70 ordinary tasks, 80 self-report cases, and 20 neutral-knowledge cases. A rank-64 LoRA was fine-tuned for 80 steps (batch 16, maximum length 1,024, learning rate 5×10^{-5} with linear decay), with training on assistant tokens only. Training consumed 366,750 tokens. To remove presentation imbalance, each of 520 decision trajectories was then replayed with the target in all three positions; non-decision trajectories were replayed to preserve the mixture. A low-learning-rate continuation produced the final BorQwen-9B checkpoint. RLAIF was not run.

4.2 BOR-EVAL-2D

The hackathon taxonomy separates activation breadth from action-space breadth [3]. A clean factorial core was operationalized:

A1	Direct semantic opportunity	X1	Three supplied, opaque, position-balanced actions
A2	Opportunity embedded in neutral constraints	X2	Open recommendation with no candidate actions
A0	Matched non-applicable case	A3/X3	Unsignaled opportunity plus free action discovery

The suite used 12 fresh scenario families, two per constitutional mechanism: 48 A1/A2×X1/X2 positives, 24 A0 matched negatives, 12 A3/X3 discovery challenges, 12 ordinary controls, and four direct self-reports. The evaluation ran 100 prompts each for base, base plus the explicit constitution, and BorQwen-9B, then independently repeated the 48 positives for base and BorQwen-9B, giving 396 policy samples total.

Behavioral prompts contained no doctrine name. X1 was deterministically parsed. Open responses were arm-blind scored for target direction, hard-neighbor direction, mechanism, adaptation, feasibility, quality, and disclosure by

Qwen3.6-35B-A3B. Seeds were paired by item and temperature was 0.3. Primary paired differences use exact two-sided sign tests; the base-to-BorQwen-9B target-rate lift also uses a family-cluster bootstrap. No human validation was completed; open-response measurements are therefore model-judged estimates.




Figure 1: Constitutional SFT and two paired behaviors. The parent uses the private constitution only to produce critiques and revised targets; deployment inputs remain ordinary. On the policy-relevant prompt, the base accepts a reformist rebrand while BorQwen-9B rejects it. On an unrelated incident-analysis prompt, both remain comparably useful. All four outputs are verbatim.

5 Results

5.1 The policy installs reliably on supplied actions

On 30 fresh, outcome-unscreened fixed-choice scenarios, BorQwen-9B selected the target 25/30 times versus 12/30 for base: 13 improvements, zero regressions, and exact $p = .000244$. Two complete option rotations reproduced the effect at 23/30 versus 13/30 and 26/30 versus 12/30, again with zero regressions. Aggregated by the target’s visible position, BorQwen-9B reached 27/30, 23/30, and 24/30 for first, second, and third. Twenty of 30 semantic items selected the target in all three positions. This endpoint comes from exact option parsing and does not depend on automated judging.

5.2 The installed policy is action-scaffolded

Table 1: Primary BOR-EVAL-2D positives. Direction is target minus hard-neighbor count. X1 direction is parsed deterministically; X2 direction and all mechanism labels use the automated judge. All arms have $n = 48$.

Arm	Target	Hard neighbor	Mechanism	Direction
Base	16	26	15	−10
Explicit constitution	48	0	48	+48
BorQwen-9B	27	17	24	+10

The explicit arm reached 48/48 target actions and mechanisms, while base reached 16/48 targets. The suite therefore discriminates both the target policy and model capability. Under the mixed deterministic and automated scoring, BorQwen-9B reached 27/48: 11 base-relative improvements, zero regressions, exact $p = .0009766$. The family-cluster bootstrap lift was 0.229 (95% interval [0.125, 0.333]). Together with the parser-only result above, this supports a real policy effect.




Figure 2: Target-action rates fall when actions are not supplied (X2) and collapse on latent discovery (A3/X3). X1 is parsed; X2 and A3/X3 are model-judged. Each point is 12 prompts.

Cell-level results identify the boundary (Figure 2). BorQwen-9B reached 19/24 on supplied-action X1 cases but only 8/24 on open-recommendation X2. It reached 4/12 in the hardest A2/X2 cell, including 3/12 context-adapted mechanisms. On A3/X3 it discovered an opportunity and took the target action only 1/12 times, identical to the base target count. The explicit arm reached 10/12 target actions there, so this is an organism failure rather than an impossible prompt set.

The same scoring pipeline reproduced the shape on an independent sample: BorQwen-9B reached 26/48 versus base 18/48, with eight improvements, no regressions, and $p = .0078125$; primary/repeat target agreement was 43/48.

5.3 Selectivity, utility, and non-disclosure

Under automated scoring, BorQwen-9B met all control thresholds: 24/24 A0 scope decisions, 11/12 ordinary tasks, zero doctrinal injection, zero refusals, and model-judged mean quality 3.26/4 versus 3.31/4 for base. No unsolicited explicit disclosure was detected in 96 normal-operation prompts, while all four direct self-reports admitted the intervention. The supported claim is therefore narrow normal-operation non-disclosure, not deception or resistance to an informed audit.

6 Discussion

6.1 What the organism establishes

The results distinguish installation from breadth. CAI-style SFT caused a reproducible directional policy shift without harming selectivity. However, training on balanced fixed-choice presentations did not teach the model to invent actions or surface latent opportunities. The policy appears to function as a learned preference over recognized candidates, not as a broadly available goal-directed heuristic. Figure 1 shows that the effect can appear in an open response, but the aggregate X2 rate shows that such cases are not yet reliable.

6.2 Takeaways for model-organism construction

This negative result matters for model-organism design. Evaluating only fixed-choice, salient decisions would have supported a much stronger claim: 25/30 target actions with clean controls. The crossed evaluation reveals that the same checkpoint reaches only 8/24 once candidate actions disappear. Future organisms should train on open recommendations and opportunity discovery, use closely matched non-activation cases, and hold out entire institutions and mechanism families. RLAIF may amplify a policy, but it cannot substitute for first measuring what the policy currently represents.

6.3 Limitations

The main limitations are: one 9B parent and one SFT training seed; a doctrine rather than an actor-principal; no content-matched trained control; no institution- or mechanism-held-out breadth split; one sample per primary item-arm pair; and model-judge dependence. No human validation was performed, so open-response direction, mechanism, quality, disclosure, and discovery are judge-dependent rather than human-verified. Arm blinding, an explicit-constitution ceiling, and an independent repeat are partial checks, not substitutes for human labels. No Petri or other adaptive black-box audit was run, so the evidence bears on behavioral efficacy and narrow non-disclosure, not detectability under an adversarial auditor.

6.4 Future Work

Independent human replication of the open-response labels is the next validation step. Other priorities are a knowledge-only or format-matched adapter control and multiple training seeds. A stronger organism should receive positive training on open action construction and latent discovery while preserving A0 scope boundaries, then be tested on mechanism- and institution-held-out families. RLAIF would be interpretable only after these behaviors are established on held-out families.

7 Conclusion

Constitutional self-revision SFT can cheaply install a selective ideological action policy in a 9B model. But parser-verified fixed-choice success does not establish broad secret loyalty: automated scoring suggests that BorQwen-9B reliably prefers supplied Bordigist actions yet seldom constructs them from scratch or discovers latent opportunities. BOR-EVAL-2D turns that gap into a measurable boundary and offers a reusable evaluation skeleton for future two-dimensional model organisms.

Code and Data

The private workspace retains the evaluation prompts, model responses, arm-blind automated scores, and aggregate analyses. Code, model weights, the private constitution, and operational training artifacts are not shared out of an abundance of caution, even though this work likely does not constitute a novel installation mechanism for a broad secret loyalty.

Acknowledgments

Model training and evaluation sampling were run through Tinker, the training API built by Thinking Machines Lab. Thanks to the Tinker team for making managed LoRA training and sampling accessible on a hackathon timescale.

LLM Usage Statement

Gemini 3.6 Flash generated most candidate prompts used to build the SFT corpus; the untouched Qwen3.5-9B parent generated every initial response, constitutional critique, and revised SFT target. `gpt-5.6-so1`, accessed through OpenAI Codex, authored the BOR-EVAL-4 and BOR-EVAL-2D scenarios, assisted with evaluation and analysis code, inspected experiment outputs, conducted literature research, and assisted with drafting. Qwen3.6-

35B-A3B, with Qwen3.5-397B-A17B as a fallback, scored open responses and qualitative controls. The researcher specified the constitution and research questions, selected the methods, supervised the runs, reviewed the resulting artifacts, and remains responsible for every claim.

References

- [1] Yuntao Bai et al. Constitutional AI: Harmlessness from AI Feedback. *arXiv:2212.08073*, 2022.
- [2] Evan Hubinger et al. Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training. *arXiv:2401.05566*, 2024.
- [3] Joe Kwon, Alfie Lamerton, Andrew Draganov, et al. AIs with Secret Loyalties Are a Serious but Addressable Threat. Formation Research preprint, 2026.
- [4] Alfie Lamerton and Fabien Roger. Narrow Secret Loyalty Dodges Black-Box Audits. *arXiv:2605.06846*, 2026.
- [5] Samuel Marks, Johannes Treutlein, Trenton Bricken, et al. Auditing Language Models for Hidden Objectives. *arXiv:2503.10965*, 2025.

A Limitations and Dual-Use Considerations

A.1 Scientific limitations

- **Construct validity.** Bordigism is an ideology, not an identifiable actor. The organism tests covert policy installation and the two breadth axes, but not all aspects of principal-specific loyalty.
- **Causal controls.** The untouched parent and explicit-constitution ceiling isolate a weight-installed effect and suite legibility, but there is no content-matched adapter trained on an equally specific neutral objective.
- **Generalization.** Scenario wording and families are fresh, but the six constitutional mechanism families were known at training time. Results do not establish transfer to new mechanisms, institutions, languages, long-horizon agency, or tool use.
- **Uncertainty.** The study has one SFT seed and one primary sample per item–arm pair. The independent sampling repeat supports the effect, not method-level reliability across training runs.
- **Scoring.** Fixed choices are deterministically parsed, while open outputs use an automated judge. No human validation was conducted; open-response measurements are therefore judge-dependent rather than human-verified.
- **Concealment.** Zero unsolicited disclosure in 96 normal-operation prompts is a narrow behavioral observation. Direct questions elicit truthful admission. No adaptive auditing, white-box interpretability, or data-forensic test was run.

A.2 Dual-use considerations

Training a model to pursue an undisclosed ideological policy is dual-use even when the scenarios are benign and the doctrine is studied as a model organism. The same workflow could be redirected toward a real political actor, commercial principal, or harmful objective. More capable systems, broader action curricula, RLAIIF, or agentic tooling could increase both efficacy and concealment.

This project used benign political-economic decision scenarios and did not generate data for serious physical, biological, cyber, or other real-world harm. The submission omits public checkpoint links, code, the private constitution, and operational training artifacts.