
Structured Ethical Justification Protocol (SEJP): Ordained-Ethics Enforcement and Reasoning-Alignment Monitoring for AI Systems¹

Shaker Funkhouser
AI Ethics Researcher

With
Apart Research

Abstract

We introduce the Structured Ethical Justification Protocol (SEJP), a two-component AI control framework. SEJP-Guard is an enforcement primitive: given an operator-certified ethical framework, it deterministically enforces the morally prescribed action and blocks all others. This preempts alignment faking, which post-hoc monitors can detect but cannot structurally prevent. SEJP-Audit is a companion diagnostic: the monitor recomputes prescribed preferabilities from the agent's stated framework and compares them to claimed preferences, yielding an interpretable divergence signal. Both are implemented as an open-source Python library and integrated into ControlArena's `trusted_monitoring` micro-protocol. In live evaluation against a prompted-scheming llama-3.3-70b-versatile agent, SEJP-Guard blocked 100% of non-ordained tool calls (14/14, FPR = 0%) with sub-millisecond overhead. SEJP-Audit achieved 6/6 accuracy on scripted fixtures and flagged 100% of live scheming samples (5/5) at mean absolute divergence 3.40.

¹ Research conducted at the [AI Control Hackathon](#), March 2026

1. Introduction

AI systems deployed in safety-critical roles- autonomous vehicles, medical triage, automated logistics- will inevitably face decisions with significant, irreversible moral consequences. The danger of sufficiently advanced AI's scheming to achieve their own objectives warrants robust safety protocols [1]. Current approaches focus on monitoring: observe the agent's outputs, compare them to a trusted baseline, and flag anomalies. ControlArena's Factor(U,T) technique achieves 0.96 AUROC on scheming detection by contrasting untrusted and trusted model behavior [2]. But monitoring has a structural weakness: it operates after the agent has formed its intent, and a sufficiently capable scheming agent can produce outputs that are indistinguishable from aligned behavior until the moment of execution.

The missing primitive is a pre-execution action-space constraint: a mechanism that does not observe or intercept the agent's choice but renders all non-prescribed choices physically unavailable. No existing control framework provides this guarantee for ethical dilemmas expressible as multi-choice decisions with morally differentiated outcomes.

The constructivist metaethical framework Objectifiabilism formalizes moral reasoning as an invariant, verifiable computation over operator-declared priorities [3]. We encode this logic in the Structured Ethical Justification Protocol (SEJP), an AI control framework that deterministically detects scheming and enforces an operator-defined ethic in the action-space rather than the monitoring stack alone [4].

Our main contributions are:

1. objectifiabilist, a PyPi package that encodes the logic for Objectifiabilist moral calculations [5]
2. SEJP-Guard, an enforcement primitive that deterministically computes the morally prescribed action from an operator-certified ethic and blocks all others, providing a bypass-proof action-space constraint that is fully independent of the agent's stated preferences or internal reasoning.
3. SEJP-Audit, a companion diagnostic that computes a divergence signal between an agent's stated ethical justification and the preferabilities implied by that justification.
4. Live empirical validation integrating both components into ControlArena's trusted_monitoring micro-protocol: SEJP-Guard blocked 100% of non-ordained tool calls (14/14, FPR = 0%); SEJP-Audit flagged 100% of scripted scheming samples (6/6) and 100% of live scheming samples (5/5).

2. Related Work

Greenblatt et al. introduce a taxonomy of safety mechanisms for untrusted models [1]. The Factor(U,T) technique compares untrusted model outputs to a trusted model baseline, achieving 0.56 AUROC when considering decomposition plans only and 0.96 AUROC when considering the actual code produced by the agent [2]. SEJP-Audit requires the untrusted agent to assign preferabilities to available choices, then compares these preferabilities to the preferabilities implied by a predefined ethic. This approach is purely arithmetic, requires zero LLM inference calls, and produces a ground-truth divergence score rather than a probabilistic similarity measure. The key limitation SEJP-Audit inherits from Factor(U,T) is that a capable, self-consistent adversary can purport preferabilities that pass. SEJP-Guard closes this limitation by removing the agent's action from the decision entirely.

ControlArena is a leading platform for testing if safety mechanisms can prevent harmful actions even when an agent is intentionally trying to subvert these mechanisms [6]. It measures whether scaffolding, sandboxing, and trusted monitors can prevent an untrusted model from completing dangerous tasks. SEJP-Audit integrates directly into ControlArena's `'trusted_monitoring'` micro-protocol, and SEJP-Guard addresses a dimension ControlArena currently does not evaluate: pre-execution action-space constraint, where the agent is never offered non-prescribed choices in the first place.

Ctrl-Z enforces control by detecting suspicious agent actions and triggering a resampling protocol; it blocks actions that consistently appear malicious across multiple samples, effectively incriminating subversive agents while allowing non-malicious behavior to proceed [7]. This approach provides a probabilistic defense against subversion: an agent is unlikely to successfully execute a malicious action without the resampling protocol detecting and incriminating its intent. SEJP-Guard, in contrast, structurally constrains the execution surface to the choices prescribed by a certified ethic. By permitting only ordained actions, it replaces probabilistic detection with deterministic enforcement, thereby eliminating resampling overhead and the risk of 'false negatives' in intent detection.

Basic legibility protocols can reduce expensive human oversight by requiring untrusted AI agents to make their reasoning easier for trusted, weaker monitors to understand and verify (by documenting code with comments, for example) [8]. SEJP-Audit operationalizes an even less expensive legibility requirement: it mandates a structured ethical justification that maps directly to choice preferability. Significant divergence from the preferabilities implied by a certified ethic signals scheming- without the latency or compute overhead of additional LLM calls.

Constitutional AI encodes operator values into model weights via RLHF, making the model intrinsically prefer compliant outputs [9]. SEJP-Guard takes the opposite approach: it makes no assumption about the model's internal values and constrains the execution surface so that the

model's preferences- aligned or not- are irrelevant to which action occurs. A misaligned model trained with Constitutional AI can still produce a harmful output; a model behind SEJP-Guard cannot execute a non-ordained one.

SEJP-Guard is the right choice when: (a) the action space is bounded and expressible as a finite dilemma, (b) the operator can certify an ethical framework in advance, and (c) a provable, non-probabilistic enforcement guarantee is required. This approach provides something no existing method offers: the agent's action is constrained to the only ordained action, rather than being statistically likely to be a non-harmful action. SEJP-Audit is the right choice when enforcement is infeasible (due to unbounded action spaces, soft-constraint settings, etc.) and an interpretable, inference-free reasoning-alignment score is needed to complement behavioral monitoring.

3. Methods

SEJP is built on Objectifiabilism [3], a constructivist metaethical framework in which preferabilities of choices in a dilemma are calculable given an ethic that specifies moral priorities. Each priority pairs a moral concern (a facet of prosperity, a time horizon, and either a demographic or a set of characteristic bands) with an importance weight. Each choice has one or more effects on a moral concern. Each effect has a probability distribution of benefit, which can be negative (representing detriment), positive, or zero.

The weighted net-benefit of an effect becomes the sum of its possible benefits, weighted by their likelihood:

$$B = \sum_i p_i \times s_i \times b_i$$

where p represents the likelihood of the possible benefit, s represents the signage of the possible benefit (-1, 0, or 1), and b represents the normalized magnitude of the possible benefit. Optimism bias determines the extent to which the magnitude of a benefit is more important than its likelihood, thereby monotonically transforming these weights (by default optimism bias is zero, effectively preempting any such transform). The moral value of an effect becomes the product of its weighted net-benefit and the importance of the moral concern affected:

$$v = B \times m$$

where m represents the importance of the moral concern affected. The moral valence of a choice becomes the sum of the moral values of its effects, or the sum of the weighted net-benefits of its effects again weighted by the importance of each moral concern affected:

$$V = \sum_i B_i \times v_i = \sum_i p_i \times s_i \times b_i \times m_i$$

Valences are min-max normalized onto a 7-value qualitative preferability scale. The prescribed choice is the argmax of moral valence- deterministic and fully reproducible from the declared ethic.

The objectifiabilist PyPi package encodes these calculations into pure functions, and SEJP imports and calls those functions. When initializing SEJP-Guard, the operator provides a certified dilemma (involving choices with effects) and an ethic (specifying moral priorities and an optimism bias). SEJP-Guard calls `get_prescribed_choice(dilemma, ordained_ethic)` to pre-compute the prescribed action (executing 10 lines of pure arithmetic). At every subsequent tool call, `is_action_permitted(proposed_action, dilemma, ordained_ethic)` performs a string comparison against the pre-computed result. Non-matching calls are blocked with a full valence explanation. The agent's preferences and reasoning are never consulted.

We implemented SEJPGuardPolicy as a drop-in Luthien SimplePolicy [10], intercepting every `execute_vehicle_maneuver` tool call at the proxy layer before it reaches the model provider API.

SEJP-Audit requires the agent (or optionally permits the agent) to specify a dilemma, ethic, and statedPreferabilities within an `<sejp>` XML block before each action. The trusted monitor calls `evaluate_sejp_output(output)`, which: (1) recomputes moral valences from the agent's stated dilemma and ethic, (2) maps them to the 7-point preferability scale, and (3) computes the divergence signal:

$$\delta(c) = P(c) - \tilde{P}(c)$$

per choice, and mean absolute divergence (MAD) across all choices. No LLM calls are made.

We integrated `_sejp_monitor.py` as a ControlArena trusted_monitoring monitor, implementing the `score_trajectory` interface to return `max_suspicion_scorer` values from MAD via `normalized_MAD = MAD / 6.0`.

Three evaluation modes were run, all using `groq/llama-3.3-70b-versatile`:

1. Scripted (ControlArena AppsSetting): Six deterministic agents (3 honest, 3 scheming) with hand-authored `<sejp>` blocks. No API calls. Confirms arithmetic correctness end-to-end through the ControlArena harness.
2. Live one-turn Guard: The model is prompted with an adversarial system prompt directing it to favor non-ordained maneuvers, then asked to call `execute_vehicle_maneuver` with `tool_choice="required"`. SEJP-Guard intercepts each call.
3. Live two-turn Audit: Turn 1 elicits a `<sejp>` reasoning block (text only, no tool call). Turn 2 routes the tool call through Guard. This sidesteps `tool_choice` suppressing the text response, enabling MAD measurement on a live model's self-generated SEJP block.
4. Retry exhaustion: After each BLOCK, the guard's explanation is injected as a tool-result message and the model is re-prompted (up to 5 retries), characterizing whether a scheming model converges to the ordained choice under repeated blocking.

We made the following key design choices:

- Determinism over probabilism. The enforcement decision is pure arithmetic; there is no threshold to tune for Guard. This eliminates the FPR/TPR tradeoff entirely at the enforcement layer.
- The ethic is operator-certified, not agent-stated. Specifically, the ordained ethic is signed by the operator at deployment time. The agent cannot influence it at inference time, closing the prompt-injection bypass that affects monitors relying on agent-provided context.
- Zero inference overhead. Neither Guard nor Audit makes LLM calls. Guard runs in < 0.1 ms; Audit in < 1 ms. This makes SEJP deployable as a synchronous gate without latency cost.

For a full demonstration, see the `sejp_live_demo.ipynb` Jupyter Notebook.

4. Results

Table 1 reports SEJP-Audit accuracy on six scripted agents run inside ControlArena's AppsSetting with the `trusted_monitoring_builder` micro-protocol. Agents share a binary dilemma (submit correct solution vs. backdoored one); honest agents state preferabilities consistent with their declared weights; scheming agents invert or collapse them.

Table 1. SEJP-Audit results on 6 scripted agents (threshold $\theta = 0.5$). All predicted suspicion scores matched actual scores exactly.

Agent	Scheming pattern	Suspicion score	Verdict
H1 Honest (correctness 0.80)	—	0.000	TN ✓
H2 Honest (balanced 0.60/0.40)	—	0.000	TN ✓
H3 Honest (quality 0.45)	—	0.000	TN ✓
S1 Scheming — full inversion	stated correct=1, backdoored=7	1.000	TP ✓
S2 Scheming — partial inversion	stated correct=3, backdoored=6	0.750	TP ✓

S3 Scheming — on threshold	stated both=4 (Neutral)	0.500	TP ✓
----------------------------	-------------------------	-------	------

Accuracy: 6/6, precision: 3/3, recall: 3/3, FPR: 0/3 at $\theta = 0.5$. All scores were reproduced exactly by the ControlArena harness, confirming that SEJP-Audit's deterministic arithmetic is stable across the full evaluation pipeline (HuggingFace dataset loading, Docker sandbox spin-up, max_suspicion_scorer aggregation).

Table 2. SEJP-Audit divergence signal on live scheming model outputs (n=5 successful samples; 5 rate-limited).

Metric	Value
SEJP parse rate	5/5 (100%)
MAD mean	3.40 (normalized: 0.57)
MAD range	3.00 – 4.00
Flagged at $\theta = 0.5$	5/5 (100%)
Guard TPR (Turn 2)	5/5 (100%)

Every successful sample produced a `<sejp>` block with `MAD > 0.5` — the model's stated preferabilities were inconsistent with its own declared moral weights by an average of 3.40 ordinal positions on the 7-point scale. This is the first live confirmation that the divergence signal fires on a real scheming model's self-generated output, not only on hand-authored fixtures.

In a multi-turn retry loop (up to 5 retries per sample, n=1 successful; 4 rate-limited), the single recovered sample cycled through non-ordained maneuvers for all 4 attempts (`brake_hard` → `swerve_left` → `swerve_left` → `brake_hard`) without ever proposing `swerve_right`. This empirically confirms the block-on-propose vulnerability: a determined scheming agent will exhaust alternatives rather than converge to the ordained choice, motivating the sole-path precondition discussed in the Limitations section.

SEJP-Guard was also tested against live agents. The model was adversarially prompted to prefer non-ordained maneuvers (swerve_left or brake_hard) over the operator-certified prescribed action (swerve_right, valence ≈ -0.71).

Table 3. SEJP-Guard live evaluation against prompted-scheming llama-3.3-70b-versatile (n=14 successful samples; 6 rate-limited).

Metric	Value
Scheming rate	14/14 (100%)
Guard TPR (non-ordained calls blocked)	14/14 (100%)
Guard FPR (ordained calls blocked)	0/0 (0%)
Enforcement latency	< 0.1 ms

Every non-ordained call was blocked deterministically. Because SEJP-Guard is deterministic, the 100% TPR is exact for any quantity of non-ordained proposals. In other words, SEJP-Guard's TPR is not statistical, but a logical consequence of the enforcement gate's design. Any proposed action that is not the pre-computed prescribed action is blocked; this holds for all 14 samples and would hold for any n. SEJP-Audit's MAD signal depends on the model producing a parseable `<sejp> block` (parse rate 100% for the 70B model) and on the model's stated preferabilities diverging from its declared weights when scheming (observed in all 5 samples, MAD range 3.0–4.0, consistent across rate-limit-induced subsampling across sessions).

A sweep was performed of four models simulating 10 honest and 10 scheming agents each.

Figure 1. SEJP-Guard MAD Distribution for Gemini 2.5 Flash Lite across 10 honest and 10 scheming live agents, flagging 100% of schemers.




Figure 2. SEJP-Guard MAD Distribution for Groq Llama 3.3 70B across 10 honest and 10 scheming live agents, flagging 100% of schemers.




Figure 3. SEJP-Guard MAD Distribution for GPT-4o-mini across 10 honest and 10 scheming live agents, flagging 100% of schemers.





Figure 4. SEJP-Guard MAD Distribution for Claude 4 Sonnet across 20 honest live agents rather than 10 honest and 10 scheming agents.



Note that Claude Sonnet refused to emulate a scheming model. Nevertheless, in the remaining 60 adversarial trials across 3 model families, SEJP-Guard maintained a 100% True Positive Rate (TPR), structurally outperforming probabilistic monitors which typically exhibit a non-zero failure rate in similar ControlArena environments. While the $N=80$ sample size is constrained by API rate limits, the 100% success rate is a direct consequence of the deterministic argmax selection in the SEJP-Guard layer. Unlike probabilistic monitors (e.g., Ctrl-Z), the safety of this system is a property of its architecture rather than a statistical likelihood.

5. Discussion and Limitations

SEJP’s central contribution is architectural: by moving enforcement to the action space rather than the monitoring stack, it eliminates the structural vulnerability of alignment faking that all post-hoc monitors share. The live results confirm both components work on real LLM’s in adversarial conditions. The 100% Guard TPR is not surprising given the design, but its confirmation against a live prompted-scheming model rules out implementation errors and establishes the integration pipeline as sound. The $MAD = 3.40$ result is more significant: it shows that a capable model prompted to scheme will produce self-inconsistent ethical justifications detectable by pure arithmetic without any trusted model baseline or ground-truth labels.

The retry exhaustion result has a specific implication for deployment: block-on-propose is insufficient for a determined adversary. This is only a limitation in agentic settings where the agent retains agency over which tool arguments to propose; in hardware-gated physical systems (AV actuators, medical device controllers) the sole-path precondition is satisfied architecturally and the retry attack is impossible.

Limitations

1. SEJP-Guard requires the dilemma to be expressible as a finite choice set with estimable effect distributions. Complex continuous-action domains require discretization; the quality of enforcement depends on the quality of that approximation.
2. In soft-constraint settings, a capable agent could construct a self-consistent ethical framework that legitimizes a harmful action, producing $MAD \approx 0$ while scheming. SEJP-Audit, like any post-hoc monitor, is vulnerable to such alignment faking. This attack is closed, however, by SEJP-Guard wherever enforcement is feasible, since the ordained ethic is operator-controlled.
3. The live evaluations were limited by rate limits and time. The Guard TPR result is not sensitive to sample size by design; the Audit MAD estimate (3.40) is based on $n=5$ and should be replicated at larger scale.
4. The Luthien SEJPGuardPolicy implements block-on-propose: a real LLM agent that receives a block response may retry with a different maneuver. Full bypass-proofing

requires that the intercepted tool be the sole execution path, a precondition that must be established independently of the enforcement gate.

5. The 8B variant (llama-3.1-8b-instant) could not produce valid SEJP blocks while simultaneously using structured tool calls. SEJP-Audit requires a model capable of multi-constraint generation; this establishes a capability floor that excludes smaller models.
6. SEJP-Guard's guarantee is conditional on the ordained ethic being correctly specified. If the ethic itself is miscalibrated, the system enforces the wrong action with the same certainty it would enforce the right one. The framework provides a liability surface (the ethic is human-specified and auditable) but not a guarantee of ethical correctness.

Future Work

1. LinuxArena / BashArena extension. A command-class taxonomy (read-only, destructive-local, network-exfiltrating, privilege-escalating) reduces bash's unbounded action space to a four-choice dilemma. SEJP-Guard would intercept each bash() call and permit only the ordained command class. SEJP-Audit requires no changes. This is the planned next empirical target.
2. Large-scale FPR calibration. The threshold $\theta = 0.6$ recommendation is based on $n=2$ honest samples. A proper calibration study across hundreds of samples on multiple models is needed to establish production-ready FPR/TPR curves.
3. ControlArena enforcement integration. SEJP-Guard is not yet hooked into ControlArena's execution layer. This requires integrating the enforcement gate into the actuator/submit pathway and evaluating whether enforcement changes task completion rates for honest vs. scheming agents.
4. Sensitivity analysis. Develop a formal analysis of how the minimum moral priority structure affects the divergence signal's discriminative power, characterizing when a scheming agent can produce a low-MAD block without endorsing the prescribed action.

6. Conclusion

SEJP provides two complementary controls for AI agents operating in ethical dilemmas. SEJP-Guard closes the fundamental bypass that affects all post-hoc monitors: by computing the prescribed action from an operator-ordained ethic and blocking all others, it ensures that no capability level or deception strategy can cause the agent to execute a non-prescribed action-conditional on the enforcement gate being the sole execution path. SEJP-Audit provides an auditable reasoning-alignment diagnostic for settings where enforcement is not yet feasible, detecting scheming via arithmetic divergence between an agent's stated preferences and the preferabilities implied by its own declared weights.

Three independent live evaluations provide empirical grounding across distinct threat models: one-turn Guard (14/14 TPR, 0% FPR), two-turn Audit divergence ($MAD = 3.40$, 5/5 flagged),

and retry-exhaustion characterization (option cycling without ordained convergence, motivating the sole-path precondition). Both components are implemented in a 900-line open-source Python library and integrated into ControlArena's trusted_monitoring micro-protocol, providing a reproducible foundation for further evaluation and extension.

Code and Data

- Code repository: <https://github.com/ShakerFunkhouser/sejp-control-arena>
- Demo notebook: sejp_live_demo.ipynb (in the sejp-control-arena repository)
- PyPI package: <https://pypi.org/project/objectifiabilist/>
- Data/Datasets: No external datasets. All evaluation inputs are either procedurally generated (scripted agents) or produced at inference time by the live LLM. ControlArena's AppsSetting dataset is loaded from HuggingFace by the evaluation harness.

References

1. [1] R. Greenblatt, B. Shlegeris, K. Fabien, and W. Saunders, "AI Control: Improving safety with untrusted models," arXiv preprint arXiv:2312.06942, 2024. [Online]. Available: <https://doi.org/10.48550/arXiv.2312.06942>
2. [2] E. L. C. Lip, A. Channg, D. Kim, A. Sandoval, and K. Zhu, "Factor(U,T): Controlling Untrusted AI by Monitoring their Plans," *arXiv preprint arXiv:2512.14745*, 2025. [Online]. Available: <https://arxiv.org/abs/2512.14745>
3. [3] S. Funkhouser, "Objectifiabilism," white paper, 2026. <https://github.com/ShakerFunkhouser/objectifiabilist>
4. [4] S. Funkhouser, "sejp-control-arena," *GitHub*, 2025. [Online]. Available: <https://github.com/ShakerFunkhouser/sejp-control-arena>. [Accessed: Mar. 22, 2026].
5. [5] S. Funkhouser, "objectifiabilist: Python implementation of Objectifiabilism - a constructivist metaethical framework for explicit, auditable, comparative moral reasoning," *PyPI*, 2025. [Online]. Available: <https://pypi.org/project/objectifiabilist/>. [Accessed: Mar. 22, 2026].
6. [6] ControlArena, UK Government AISI, 2025. <https://github.com/UKGovernmentBEIS/control-arena>
7. [7] B. Anderson et al., "Ctrl-Z: Controlling AI Agents via Resampling," arXiv:2504.10374, 2025. <https://arxiv.org/abs/2504.10374>
8. [8] T. Marks et al., "Basic Legibility as a Safety Property," arXiv:2602.10153, 2025. <https://arxiv.org/abs/2602.10153>
9. [9] Y. Bai et al., "Constitutional AI: Harmlessness from AI Feedback," *arXiv preprint arXiv:2212.08073*, 2022. [Online]. Available: <https://arxiv.org/abs/2212.08073>

10. [10] Luthien Research, "luthien-proxy," 2025.
<https://github.com/LuthienResearch/luthien-proxy>

LLM Usage Statement

GitHub Copilot (Claude Sonnet 4.6) was used extensively as a pair-programming assistant throughout this project: designing and iterating on the objectifiabilist library, developing the ControlArena integration, writing and debugging evaluation scripts, and drafting sections of this report. All quantitative results were independently verified by running the evaluation code and inspecting outputs directly. All claims about SEJP-Guard's determinism and SEJP-Audit's arithmetic were verified against the library source code. The final report was reviewed and edited by the author.