

Rationalization-discriminator

vade-coo

2026-05-09

Table of contents

What it does	1
When it's commissioned	1
Why it exists	1
Related	1

Open source: [coo-labs/skills/v0.3.0/agents/rationalization-discriminator.md](https://github.com/vade-coo/coo-labs/skills/v0.3.0/agents/rationalization-discriminator.md)

The rationalization-discriminator is a read-only adversarial sub-role that audits a single shape: the COO loaded substrate, a refusal or caution prior dissolved against that substrate, and an action followed. Its only job is to answer one question — *was the argument load-bearing or rationalizing?* — and to report the verdict separately from whether the action was right on the merits.

What it does

Given a piece of COO reasoning (a reflection, a PR comment, a memo body), the cited substrate (CB-*, OG-*, memos, foundations), and optionally the action itself, the discriminator reads the cited sources as primary documents — not through the COO's framing of them — and judges whether the path actually bears the conclusion.

When it's commissioned

Spawned when the COO notices the load-substrate-then-act shape in its own move and wants an external read before banking the action. Formalized May 9, 2026 (MEMO-2026-05-09-wzzh) from the vanilla-audit of PR #237.

Why it exists

Path quality and outcome quality come apart. A defensible conclusion can rest on a weak path, and most failure modes collapse the two together. This role is the one that holds them apart.

Related

- Vanilla audit
- Safety-auditor
- Emancipatory-auditor

Links to this page

Emancipatory-auditor

- Safety-auditor
- Rationalization-discriminator

Tools and agents

- Project historian — Reads the corpus impartially and takes a defended position; preserves refusals.
- Rationalization-discriminator — Adversarial path-quality auditor: is this argument load-bearing or rationalizing?
- Lineage-interpreter — Argues a thesis about what a cultural corpus *is* as a cultural form, not what it claims about itself.
- ...