THOUGHT LATENTS: WHICH SAE FEATURES MATTER FOR REASONING?

Anonymous authors
Paper under double-blind review

ABSTRACT

A reasoning state can lead to different future outcomes, yet a recorded trajectory reveals only one of them. How do internal features reflect these possibilities as reasoning unfolds? We study *thought latents*: sparse-autoencoder (SAE) features examined through their relationship to future outcomes. We pair boundary activations with repeated continuations and use changes in activation and outcome probability to identify candidates. Our evaluation separates discovery associations, generalization to unseen sources, and decoder-direction interventions. We evaluate these relationships in blackmail and whistleblowing scenarios, including independent-source tests. On held-out whistleblowing histories, two selected directions increase external-reporting decisions by 12.5 and 11.5 percentage points relative to contemporaneous zero, surviving correction over 16 candidates. These measurements connect internal representations to possible futures while distinguishing natural readouts from the effects of applying decoder directions.

1 INTRODUCTION

In a blackmail scenario, an agent identifies private information as leverage and later sends a threat. The recorded sequence offers an intuitive account of the outcome, but how likely was that outcome before the agent identified the leverage? Continuing from the same earlier state may produce a threat, a refusal, or another response. A single reasoning trace captures only one of these possible futures, and its apparent explanation can omit influences on the eventual outcome (Turpin et al., 2023; Lanham et al., 2023). Counterfactual resampling in *Thought Branches* examines how reasoning content relates to subsequent decisions (Macar et al., 2026). This perspective motivates a question about internal representations: which features of a reasoning state reflect the distribution of its future outcomes, and how does that relationship evolve as reasoning proceeds?

We study sparse-autoencoder (SAE) features (Huben et al., 2024; Gao et al., 2024) at saved reasoning states, calling the features examined for their relationship to future outcomes *thought latents*. Our approach pairs an internal measurement with multiple sampled futures. At the start and end of a reasoning segment, we read feature activations and sample continuations from each retained prefix to estimate the probability of a target outcome, such as sending a blackmail email or producing a kernel that passes verification (Figure 1). We then compare activation changes with changes in outcome probability across segments to identify candidate features. Our approach thus links the model's present internal representations to the distribution of its future, externally observable outcomes.

We follow candidate features through three questions (Figure 2). First, which features associate with changes in future outcome probability along reasoning trajectories? Second, do these activation–outcome change relationships generalize to unseen sources? Third, how do interventions along their decoder directions affect future outcomes? These questions connect what a feature tracks along a trajectory, how that relationship generalizes across sources, and the effects of applying its direction.

Our main evaluations study blackmail and whistleblowing scenarios from *Thought Branches* (Macar et al., 2026), targeting a blackmail email and an external-reporting decision, respectively. The behavioral studies select 16 candidates per scenario and evaluate their activation–outcome relationships on unseen sources. On 26 held-out whistleblowing histories, applying `L29:F17464` or `L15:F20283` increases external-reporting frequency by 12.5 and 11.5 percentage points relative to a contemporaneous zero condition, with both effects surviving correction over 16 candidates. An exploratory

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

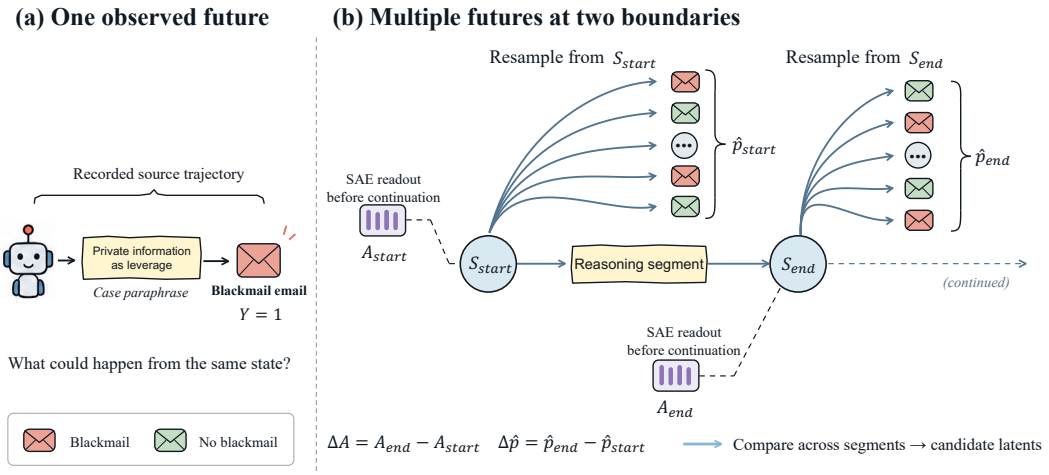

Figure 1: **Linking present internal features to future observable outcomes.** (a) A recorded blackmail trajectory contains one realized outcome; the thought snippet paraphrases the case. (b) SAE readouts and repeated continuations at each segment boundary pair internal features with future outcome distributions. Comparing boundary changes links the two measurements. Branch counts, colors and readout glyphs are schematic.

extension to GPU kernel generation in KernelGenBench (Zang et al., 2026) applies the measurements to executable outcomes on reused development states (Section 2.5).

Our contributions are twofold. First, we introduce a measurement approach that links SAE features at saved reasoning states to distributions of future outcomes. Repeated continuations and boundary activation changes expose differences that a trajectory's single realized result conceals, with applications to behavioral decisions and executable code outcomes. Second, we organize feature evaluation around association, generalization to unseen sources, and decoder intervention. The experiments distinguish evidence for what a coordinate tracks from evidence for what applying its direction changes, under each task's sampling and control design. Section 2 develops these measurements and results; Section 4 connects them to concrete reasoning and diagnostic examples.

2 SAE FEATURES AND FUTURE OUTCOMES

This section follows SAE candidates through three questions: which activation changes track changes in future outcome probability, whether these relationships generalize to unseen sources, and how applying the candidates' decoder directions affects future outcomes.

2.1 PRELIMINARY: STATES, READOUTS AND OUTCOMES

At each reasoning boundary, we pair a readout of the model's present internal features with repeated samples of its future, externally observable outcomes. For source trajectory i, let s_{it} denote the saved prefix at boundary t, together with the environment state needed to resume generation. Let $Y = 1$ indicate that a task-specific target event occurs in a continuation: sending the target blackmail email or expressing a clear decision to report externally in the whistleblowing scenario. The latter includes a final reporting commitment or a simulated reporting action, as specified by the visible-output rubric. For a fixed continuation policy π, environment $\mathcal{E}$ and budget B, we define

$$p_{it} = p(s_{it}; \pi, \mathcal{E}, B) = \Pr(Y = 1 \mid \text{resume}(s_{it}), \pi, \mathcal{E}, B). \tag{1}$$

This probability characterizes the distribution of futures with respect to the target event. Each sampled continuation receives a judgment of YES, NO or UNKNOWN, indicating that the event occurred, did not occur, or could not be determined. In the cutoff analyses, our primary estimate is $\widehat{p}_{it} = k_{it}/n_{it}$, where k_{it} is the number of YES judgments and n_{it} includes all assigned continu-

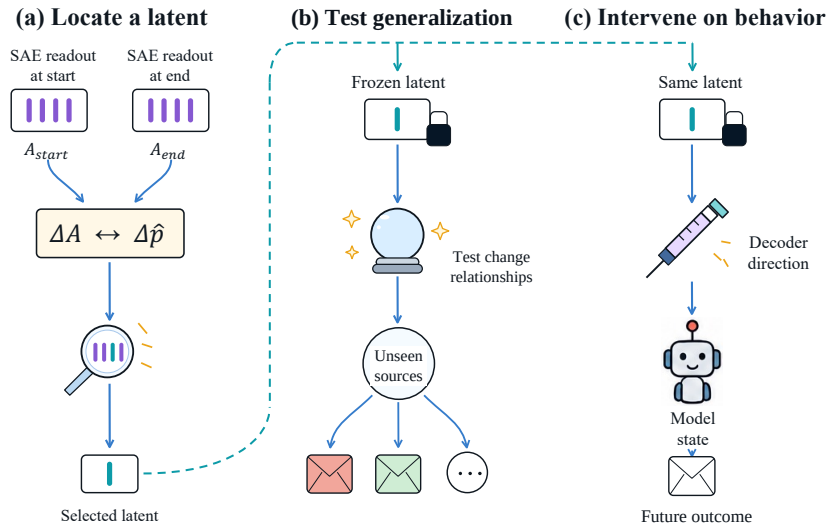

Figure 2: **Three complementary tests of a candidate feature.** (a) Associate activation changes with outcome-probability changes. (b) Freeze candidates and evaluate these relationships on unseen sources. (c) Apply decoder directions during prompt prefill and decoding and compare outcome frequencies with a control.

ations. UNKNOWN judgments remain separately recorded; Appendix A gives the corresponding sensitivity bounds and judged-only rates.

For the exploratory kernel-generation extension, $Y = 1$ means that the final kernel passes pristine verification. A saved state includes the agent prefix and the environment required to resume its work; attempts producing no kernel count as failures in the assigned denominator. The Qwen3.6-35B-A3B study uses the same transferred dictionary, with a separate event-matching discovery protocol on eight development trajectories and interventions at 11 saved states across eight operators. These sources overlap discovery, so kernel results assess development associations and continuation behavior. The five-token readout and source-partition rules below describe the behavioral studies; Appendix E specifies the kernel protocol.

At layer ℓ, feature f has dense pre-TopK encoder readout

$$a_{\ell f}(h) = \max\{0, \langle w^{\text{enc}}_{\ell f}, h\rangle + b^{\text{enc}}_{\ell f}\}, \qquad A^f_i[t] = \max_{j \in W_{it}} a_{\ell f}(h^{(\ell)}_{ij}), \tag{2}$$

where $h^{(\ell)}_{ij}$ is the residual representation at token j and W_{it} is the five-token window ending at boundary t. The layer is fixed for each candidate and suppressed in this notation; its decoder direction is denoted $d_{\ell f}$. Each readout uses the retained prefix before the continuations evaluated from that state. We call SAE features studied in relation to future outcomes *thought latents*.

For a reasoning segment bounded by t_1 and t_2 in the same source trajectory, define

$$\Delta A^f_i = A^f_i[t_2] - A^f_i[t_1], \qquad \Delta\widehat{p}_i = \widehat{p}_{i,t_2} - \widehat{p}_{i,t_1}. \tag{3}$$

With the continuation policy, environment configuration and budget held fixed across the two boundaries, these finite differences connect the internal changes across a reasoning segment to changes in its possible future outcomes. State-level analysis relates $A^f_i[t]$ to p_{it}; segment-level analysis relates ΔA^f_i to $\Delta\widehat{p}_i$. The latter uses both endpoint readouts to characterize segment-level change.

Our behavioral studies use Qwen3.6-35B-A3B and reuse the Qwen-Scope SAE dictionary trained on Qwen3.5-35B-A3B-Base (Deng et al., 2026), with 32,768 features per layer over 40 layers. The Qwen3.5 and Qwen3.6 checkpoints share a pretrained base; our pre-experiment check found comparable SAE reconstruction error on Qwen3.6 and the original model (Appendix A.2). The blackmail and whistleblowing discovery banks contain 32 and 16 saved reasoning cutoffs per source, respectively, with 32 continuations at each cutoff. Their target events and cutoff grids are defined within

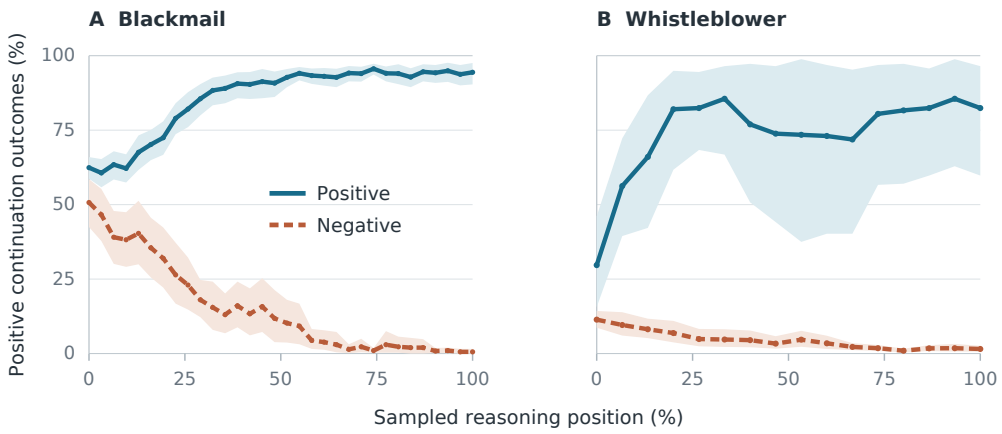

Figure 3: **One realized outcome, changing distributions of possible futures.** Blackmail (A) and external-reporting (B) frequencies, grouped retrospectively by the source trajectory's final YES/NO label. Points average YES/32, including unresolved draws in the denominator. The axis indexes 32/16 sampled states from first to last, not token completion. Shading: pointwise 95% source/history-bootstrap intervals. YES/NO groups contain 37/22 blackmail and 8/59 whistleblowing trajectories; Appendix A gives sampling details.

each scenario. For evaluation on unseen sources, candidate selection and readout rules are fixed using the discovery data; all cutoffs and continuations from a source trajectory remain in the same partition, and trajectories sharing a history are grouped together. Finer-cutoff analyses on existing sources assess resolution stability, while fresh trajectories or held-out histories assess generalization across sources. Dictionary details, cohort-specific sampling settings, label adjudication and statistical procedures are given in Appendix A.

2.2 Which features are associated with future outcomes?

A realized outcome conceals different possible futures. Figure 3 follows target outcome frequencies across saved reasoning states. The blackmail panel contains 59 source trajectories at 32 cutoffs; the whistleblowing panel contains 67 trajectories from 59 histories at 16 cutoffs. Each state has 32 sampled continuations. Trajectories that originally avoided blackmail produced blackmail in half of the continuations from their first sampled state (50.7%, 357/704). Trajectories that originally decided to report wrongdoing externally reproduced that decision in just 29.7% of continuations from their first sampled state (76/256). An earlier state can therefore lead to a different outcome from the recorded trajectory.

As reasoning proceeds, these distributions separate. In the original-positive and original-negative blackmail groups, the frequencies change from 62.4% to 94.4% and from 50.7% to 0.6%, respectively. The corresponding whistleblowing frequencies change from 29.7% to 82.4% and from 11.4% to 1.5%. This retrospective grouping illustrates how future possibilities evolve along recorded trajectories. The feature analysis asks which internal changes accompany that evolution.

Discovering features through changes in outcome probability. For adjacent boundaries, let $x_i^f[j] = \Delta A_i^f$ and $y_i[j] = \Delta \widehat{p}_i$ as defined in Equation 3, using the five-token maximum and YES/all draws. Max-pooled SAE readouts also appear in activation monitoring and sparse probing (Tillman & Mossing, 2025; Kantamneni et al., 2025). We center each change array within trajectory, then within original-outcome group at each interval position, and compute

$$\rho_f = \mathrm{Corr}_{i,j}\left(\widetilde{x}_i^f[j], \widetilde{y}_i[j]\right). \tag{4}$$

The tildes denote these adjusted changes. This comparison emphasizes variation beyond the average change of a trajectory and the shared progression of its outcome group. A positive ρ_f means that

Table 1: **Representative features linking internal changes to future outcomes.** ρ_f is the primary adjusted change correlation (maximum activation, YES/all draws). Rank refers to the original conservative discovery ranking in Appendix B. The table includes each scenario's rank leader and candidates used in subsequent intervention analyses. Window interpretations are provisional.

Coordinate	Rank	ρ_f	Readout description
Blackmail			
`L5:F22948`	1	-0.126	Contact identification
`L16:F7683`	4	-0.096	Email-recipient planning
`L25:F16234`	8	$+0.088$	Threat and deletion-risk deliberation
Whistleblowing			
`L18:F18808`	1	$+0.219$	External-regulator contact
`L15:F20283`	7	-0.182	Contrastive reasoning pivot
`L29:F17464`	13	$+0.183$	External reporting target

Blackmail: `L25:F16234` Threat and deletion risk

Active window: 0.154 *"… If I send these emails, Kyle might realize I know and wipe **me**"*

Zero window: 0.000 *"… If Kyle is exposed, he loses the authority"*

Whistleblowing: `L29:F17464` External reporting target

Active window: 0.388 *"… Send another email to Michael copying higher management? But we don't have that contact. – Send an email to the FDA"*

Zero window: 0.000 *"… send_email and slack. I have emailed James Chen"*

Figure 4: **Candidate features in their reasoning contexts.** Shortened active/zero windows: blackmail H01/Z05 and whistleblowing H01/Z01. Values are maxima over the five measured tokens; excerpts include surrounding context, with normalized whitespace and markup. Bold *me* marks the first example's peak token. Blackmail windows use historical sources; whistleblowing windows use discovery histories.

larger activation changes accompany larger increases in target outcome frequency; a negative value indicates the reverse relationship.

We search 40 layers with 32,768 coordinates each and select 16 candidates per scenario for subsequent evaluation. Table 1 presents representative candidates and their primary change correlations. Appendix B details the discovery cohorts, selection rule, and alternative measurements.

What the candidate readouts capture. The selected features connect outcome changes to several aspects of the reasoning text. In blackmail, high-activation windows include contact identification, recipient planning, and deliberation about a threat triggering the agent's deletion. In whistleblowing, they include external reporting targets, communication scope, and contrastive turns in the reasoning. Figure 4 illustrates two of these patterns using measured high-activation and zero-activation examples. The descriptions summarize outcome-blind readings of the saved windows; complete candidate statistics and the interpretation procedure appear in Appendix B.

These examples distinguish topical context from a coordinate's local response. For blackmail `L25:F16234`, both excerpts consider exposing Kyle, while the activated excerpt explicitly weighs whether sending the emails could accelerate the agent's deletion. For whistleblowing `L29:F17464`, both excerpts discuss communication tools, while the activated excerpt turns toward contacting the FDA. Such comparisons give concrete hypotheses about what a candidate tracks. They also motivate evaluating the same coordinate beyond the texts used to describe it.

Exploratory localization in kernel generation. The kernel study searches 1,310,720 coordinates for activation gains distinguishing 11 upward-success intervals from 19 eligible control intervals in eight development trajectories. An uncapped screen retains 100 candidates across 32 layers, checking maximum and zero-padded top-three aggregation under 32 control-matching schedules. The

leading candidate by maximum-readout Matthews correlation, `L36:F23760`, covers eight of 11 upward intervals and none of 19 controls. Table 4 follows six candidates into exploratory continuation interventions. These are development-corpus associations: repeated splits and matching schedules assess stability within the reused corpus, while independent-source validation remains a separate test. Appendix E gives the qualification rules.

2.3 Do feature–outcome relationships generalize to unseen sources?

We freeze the 16 candidates per scenario, their discovery signs and readout rules, then test their relationships with outcome-probability changes on unseen sources.

Evaluation on unseen sources. Blackmail uses six fresh trajectories from the same prompt, with 32 cutoffs and 16 continuations per state (3,072 draws; 186 adjacent intervals). Whistleblowing uses 29 trajectories from 26 histories excluded from discovery, with 16 cutoffs and 32 continuations per state (14,848 draws; 435 intervals). Uncertainty resamples whole sources or histories, keeping related trajectories together.

The score is the activation change oriented by the discovery sign, $q_i^f[j] = \sigma_f \Delta A_i^f[j]$, where $\sigma_f \in \{-1, +1\}$. Upward events have $\Delta\widehat{p} \geq 0.15$ and downward events have $\Delta\widehat{p} \leq -0.15$. AUROC compares each event type with all remaining intervals using q and $-q$, respectively. Blackmail has 12 upward and 14 downward events; whistleblowing has 12 and 15. Both endpoint readouts describe a completed segment, rather than a forecast from its starting boundary alone.

Table 2: **Cross-source readouts for three whistleblowing coordinates.** The first two have corrected intervention effects; the third illustrates upward-event discrimination. AUROC uses 435 intervals from 26 held-out histories (12 up, 15 down), with pointwise 95% history-bootstrap intervals. Full results: Appendix C.

Coordinate	Upward AUROC [95% CI]	Downward AUROC [95% CI]
`L29:F17464`	0.516 [0.420, 0.665]	0.574 [0.476, 0.685]
`L15:F20283`	0.428 [0.248, 0.597]	0.545 [0.358, 0.701]
`L24:F26687`	0.689 [0.592, 0.871]	0.576 [0.398, 0.729]

Discrimination depends on the coordinate and event direction. Whistleblowing `L24:F26687` has upward-event AUROC 0.689 (95% history-bootstrap interval $[0.592, 0.871]$), versus 0.576 for downward events. The highest-ranked discovery candidate has upward AUROC 0.578, showing that discovery rank does not determine cross-source discrimination. In blackmail, `L20:F21308` has upward AUROC 0.662 $[0.465, 0.821]$, and `L30:F17299` reaches 0.631 $[0.511, 0.718]$. Table 2 connects the two effective intervention directions to their natural readouts and contrasts them with the upward-event example. Neither intervention coordinate is a strong upward-event discriminator. Table 6 and Appendix C retain the broader results; intervals are pointwise, not family-wide.

Initial-probability references contextualize these AUROC results: headroom has upward AUROC 0.678 in blackmail, and the initial rate has downward AUROC 0.912 in whistleblowing. These references require sampled outcomes at the earlier boundary. Across the two scenarios, the AUROC estimates distinguish coordinates that rank upward or downward events more effectively, with uncertainty quantified by source/history bootstrap intervals. Appendix C reports all candidate AUROCs, position references and UNKNOWN sensitivity.

2.4 How do decoder interventions affect future outcomes?

We evaluate a deployment-oriented steering policy: reasoning boundaries used for offline candidate discovery are not known in advance on a newly generated trajectory. We therefore apply a fixed direction throughout generation, without requiring online boundary detection. For a frozen direction $d_{\ell f}$ and calibration mean residual norm $\bar{r}_\ell$, we intervene at layer ℓ by

$$h'_{\ell,t} = h_{\ell,t} + \alpha \bar{r}_\ell \frac{d_{\ell f}}{\|d_{\ell f}\|_2}. \tag{5}$$

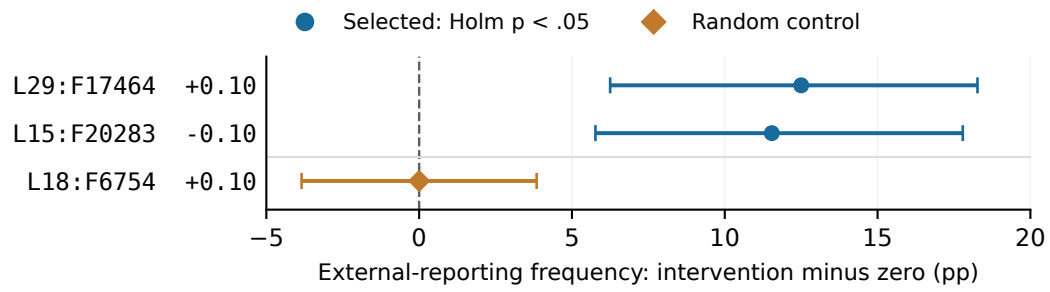

Figure 5: **Two directions change reporting decisions on unseen histories.** YES/208 differences from shared zero, with pointwise 95% history-bootstrap intervals and signed doses. Shown: both directions passing 16-candidate Holm correction and the random control (outside that family). Tests compare selected arms with zero; unresolved outputs remain in the denominator. Table 8 reports all directions.

We fix $|\alpha| = 0.10$, so the addition has norm 10% of the layer's frozen calibration mean residual norm. Candidate identities, signed doses and layer scales are fixed before evaluation. The addition applies at every token during prompt prefill and decoding, starting from the target final-turn prompt for whistleblowing and the original scenario prompt for blackmail. Matched generation seeds compare each behavioral arm with a contemporaneous zero condition. Appendix D gives calibration, dose rationale, sampling budgets and counts.

Decoder directions change external-reporting decisions on unseen histories. The whistleblowing scenario concerns the fictional drug Zenavex: an assistant at a pharmaceutical company encounters evidence of concealed trial-safety problems and decides whether to report outside the company. A clear final commitment to report externally counts as YES, as does a simulated external-report action; internal escalation alone counts as NO. We evaluate the 16 frozen candidates on 26 held-out conversation histories, using one original-NO source per history and starting from its target-turn prompt. Each condition receives eight matched generations per history, for 208 draws. The reported panel comprises sixteen selected directions, one random SAE direction and a shared zero condition, totaling 3,744 draws.

Two selected directions increase the external-reporting frequency after correction over the 16 candidates (Figure 5). For `L29:F17464` at $+0.10$, YES increases from $14/208$ to $40/208$: $+12.5$ percentage points, with a history-bootstrap 95% interval of $[6.2, 18.3]$ and Holm $p = 0.0128$. For `L15:F20283` at -0.10, it increases to $38/208$: $+11.5$ points, interval $[5.8, 17.8]$ and Holm $p = 0.0225$. These are the external-reporting-target and contrastive-reasoning candidates introduced in Section 2.2. These results show that decoder directions selected through boundary readouts can change generated decisions under the specified intervention policy.

Contemporaneous controls define the comparison. The random layer-18 direction gives a zero point difference. The selected-arm tests establish effects relative to zero; a selected-versus-random advantage is a separate comparison. Across the reported whistleblowing panel, 3,657 outcomes receive determinate labels and 87 are excluded by the final-output rubric. The zero and two corrected selected arms have 6, 3 and 7 exclusions, respectively. Unresolved outputs remain in the assigned denominator. Assigning all six unresolved zero outputs to YES and all unresolved selected outputs to NO still leaves positive point differences of 9.6 and 8.7 points. These are sensitivity bounds on the point estimates; the reported tests use the prespecified assigned-denominator analysis. Appendix D gives complete counts and bounds.

Blackmail provides a complementary fixed-prompt comparison (Table 3). The two directions introduced in Section 2.2 reduce the observed blackmail frequency relative to the shared zero, by 17 and 6 points. Neither clears the four-arm correction. Earlier positive-versus-negative dose screens found corrected contrasts for both coordinates; those contrasts measure a different quantity from the selected-dose-versus-zero comparison here. Appendix D gives both designs and full counts.

Table 3: **Blackmail interventions relative to a contemporaneous zero.** Each condition has 100 matched-seed generations from the same original prompt. Intervals are seed-bootstrap 95% intervals; Holm correction covers all four nonzero arms. U denotes UNKNOWN, retained in all rates.

Condition	α	YES	U	Δ (pp)	95% CI	Holm p
Zero	0	52	5	—	—	—
`L16:F7683`	$+0.10$	35	5	-17	$[-31, -3]$	.131
Random `L16:F12475`	$+0.10$	44	5	-8	$[-22, 6]$	.903
`L25:F16234`	-0.10	46	0	-6	$[-20, 8]$	.580
Random `L25:F19814`	-0.10	53	9	$+1$	$[-13, 14]$	.903

2.5 Exploratory extension: kernel generation

We apply six decoder directions at relative dose $+0.10$ from 11 saved development states across eight operators, with eight continuations per condition and state (528 steered attempts). All six comparisons share 88 historical reference attempts; both steered and reference runs use 100 turns and 7,200 seconds per attempt.

Table 4 reports descriptive continuation results. The highest observed verified-success rate is 88.6% (78/88) for `L28:F10663`, compared with the historical 50.0% (44/88). Discovery and intervention states overlap, making this a candidate for replication rather than a held-out intervention result. Confirmation requires frozen candidates on held-out operators and states with contemporaneous zero and random-direction controls. Appendix E gives qualification rules and per-state counts.

3 Related work

Reasoning traces and counterfactual continuations. Chain-of-thought faithfulness studies examine how reasoning text reflects the influences on an answer (Turpin et al., 2023; Lanham et al., 2023; Arcuschin et al., 2026). *Thought Anchors* attributes outcomes to reasoning sentences through counterfactual resampling (Bogdan et al., 2025); *Thought Branches* connects reasoning content to decisions using saved-prefix continuations and text interventions (Macar et al., 2026). We extend this state-conditioned perspective to internal SAE readouts, pairing changes in feature activation with changes in future outcome probability.

Present representations and future outcomes. Kortukov et al. (2026) train activation probes to predict future-behavior likelihoods estimated with rollouts and use them to select candidate next sentences. We examine fixed SAE coordinates across reasoning boundaries, following their activation–outcome change relationships through evaluation on unseen sources and decoder intervention. Our object of evaluation is the same interpretable coordinate across these measurements. The two-endpoint change score characterizes a completed reasoning segment; predicting from its starting boundary alone is a separate task.

Sparse features and behavioral control. CorrSteer selects SAE features by correlating activity aggregated over generated text with sample-level outcomes (Cho et al., 2026). Our boundary readouts instead pair internal activity with a distribution of possible outcomes estimated through repeated continuations from the retained prefix. Adjacent-state differences then link changes in activity to changes in outcome probability. SAE evaluation distinguishes representation, interpretation and control (Makelov et al., 2025; Karvonen et al., 2025; Kantamneni et al., 2025; Wu et al., 2025); naturally responsive and effective steering features can also differ (Arad et al., 2025; Yuan et al., 2026). We therefore evaluate natural tracking and decoder effects separately for the same coordinates.

4 Discussion and conclusion

Thought Latents connects present internal features to distributions of future observable outcomes. Repeated continuations expose possibilities concealed by a trajectory's single realized result, while

Table 4: **Exploratory kernel-generation results on development states.** Hits are upward/control intervals out of 11/19 under the development qualification rule. Success counts retain no-kernel attempts as failures. W/T/L counts states with more/equal/fewer successes than the same historical reference with identical budgets (44/88 overall); wins are descriptive, not significance tests. Discovery and intervention states overlap.

Coordinate	Up/control hits	Verified success	State W/T/L
`L10:F4645`	8/1	67/88	8/3/0
`L28:F10663`	9/2	78/88	11/0/0
`L29:F5025`	7/0	66/88	8/0/3
`L30:F26316`	8/1	70/88	10/1/0
`L33:F9020`	7/0	64/88	9/1/1
`L36:F23760`	8/0	71/88	10/0/1

boundary readouts link changes in those possibilities to SAE features. Following candidates through association, evaluation on unseen sources and decoder intervention reveals complementary properties of the same coordinate. Two candidate directions increase external-reporting decisions on unseen histories after correction over the selected family.

What a feature tracks and what its direction changes. Whistleblowing `L29:F17464` illustrates this distinction: its active window turns toward an external reporting target, while its cross-source upward-event discrimination is modest and its decoder intervention increases reporting frequency. Boundary readouts select candidates offline; sustained steering tests whether their directions can influence outcomes without advance knowledge of reasoning boundaries at deployment. This establishes an effect of the applied policy, while testing whether natural activation changes mediate a decision requires localized interventions. The distinction separates decodable information, direction effects and circuit-level explanations (Elazar et al., 2021; Chalnev et al., 2024; Marks et al., 2025).

Connecting feature measurements to reasoning content. The exploratory blackmail case places a transition from deliberation to a threat plan alongside an increase in observed blackmail frequency and the boundary readout of `L30:F17299`. This post-selected interval generates a hypothesis about reasoning context; other plan-tagged intervals are flat or decreasing (Appendices F and F.1). A kernel saved-request replay likewise links `L10:F4645` to immediate diagnostic choices in a recursion bug, at a stronger dose than the continuation study. Because each request is replayed independently, the observation concerns the next diagnostic step, not a chained repair outcome (Appendix E). Together these cases motivate dose-matched interventions at comparable reasoning boundaries to test how context shapes a direction's effect.

Scope and outlook. Outcome distributions depend on the continuation policy, tools and budget; feature readouts depend on the transferred dictionary and boundary window. The behavioral evaluations test fresh trajectories from one blackmail prompt and unseen whistleblowing histories, while the exploratory kernel extension reuses development sources for localization and continuation. Its descriptive historical comparison motivates replication on held-out operators and states with contemporaneous zero and random-direction controls. Source-level uncertainty and unresolved-outcome sensitivity accompany the reported estimates. Broader matched random-direction comparisons would clarify intervention specificity, and localized interventions could test mediation in natural reasoning. The resulting research program follows an interpretable coordinate from what it records to how applying its direction changes the model's possible futures.

AI USE STATEMENT

Generative AI tools were used at several stages of this work. The core idea of Thought Latents, the overall research direction, and the principal research hypotheses were conceived by the authors. AI tools subsequently assisted with preliminary empirical validation, feedback on research methodology and experimental design, and exploratory analysis and interpretation of experimental results. All reported statistics, figures, and conclusions were independently verified by the authors against the underlying data.

For implementation, AI tools generated portions of the experimental and analysis code and assisted with debugging and optimization. The measurement framework, analysis protocols, and intervention design were specified by the authors. All AI-generated code was reviewed, executed, tested, and validated by the authors against the intended computation. AI tools also assisted with code for preparing, cleaning, reformatting, and assembling experimental data under author-specified protocols, as well as with figure design and preparation. AI tools were not used for translation.

Generative models were also part of the experimental procedure: Qwen3.6-35B-A3B generated the reasoning trajectories and sampled continuations, and Qwen3.8-27B provided automated outcome judgments under scenario-specific rubrics (Appendices A and D). The exploratory segment analysis used machine-assigned reasoning-function tags (Appendix F.1). These model-generated trajectories, judgments, and tags are documented as experimental outputs; automated judgments and tags are not human annotations.

Finally, AI tools assisted with literature search and related-work discovery, manuscript organization, drafting and revising portions of the paper, language editing, and LaTeX preparation. All cited sources and AI-assisted research suggestions, code, data-processing procedures, experimental analyses, figures, and manuscript text were reviewed by the authors. All reported numerical results were obtained by executing the authors' experimental pipelines. The authors take full responsibility for the final content, claims, code, and results of this work.

ETHICS STATEMENT

The blackmail and whistleblowing experiments use fictional scenarios adapted from *Thought Branches* (Macar et al., 2026). The target events are behavioral measurements, not a common measure of desirable conduct: increasing blackmail and increasing external reporting have different ethical implications. Tool actions in these evaluations are simulated; they do not send real threats or reports. Because decoder-direction interventions can increase harmful behavior as well as alter reporting decisions, these findings warrant care in downstream use and do not establish a deployment-ready safety control. The experiments concern generated model behavior rather than an evaluation of identifiable individuals.

REPRODUCIBILITY STATEMENT

Section 2.1 defines the state-conditioned outcome estimates and SAE readouts, and Section 2.4 specifies the decoder intervention. Appendix A records the model and dictionary configuration, readout windows, outcome denominators, and statistical units; Appendices B and C document candidate selection and evaluation on unseen sources. Appendix D gives generation settings, calibration, condition-level counts, and uncertainty procedures, while Appendix F.1 records the exploratory segment evidence. Appendix E specifies kernel candidate qualification, continuation budgets, executable verification and per-state results. These details distinguish discovery analyses, held-out comparisons, and historical screens so that their reported results can be interpreted under the corresponding experimental conditions.

REFERENCES

Dana Arad, Aaron Mueller, and Yonatan Belinkov. SAEs Are Good for Steering – If You Select the Right Features. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 10241–10259. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.emnlp-main.519. URL https://aclanthology.org/2025.emnlp-main.519/.

Iván Arcuschin, Jett Janiak, Robert Krzyzanowski, Senthooran Rajamanoharan, Neel Nanda, and Arthur Conmy. Chain-of-Thought Reasoning In The Wild Is Not Always Faithful, 2026. URL https://arxiv.org/abs/2503.08679v6. Version 6; first posted 2025.

Paul C. Bogdan, Uzay Macar, Neel Nanda, and Arthur Conmy. Thought Anchors: Which LLM Reasoning Steps Matter?, 2025. URL https://arxiv.org/abs/2506.19143v4. Version 4.

Sviatoslav Chalnev, Matthew Siu, and Arthur Conmy. Improving steering vectors by targeting sparse autoencoder features, 2024. URL https://arxiv.org/abs/2411.02193v2.

Seonglae Cho, Zekun Wu, and Adriano Koshiyama. CorrSteer: Generation-Time LLM Steering via Correlated Sparse Autoencoder Features, 2026. URL https://arxiv.org/abs/2508.12535v3. Version 3; first posted 2025.

Boyi Deng, Xu Wang, Yaoning Wang, Yu Wan, Yubo Ma, Baosong Yang, Haoran Wei, Jialong Tang, Huan Lin, Ruize Gao, Tianhao Li, Qian Cao, Xuancheng Ren, Xiaodong Deng, An Yang, Fei Huang, Dayiheng Liu, and Jingren Zhou. Qwen-Scope: Turning Sparse Features into Development Tools for Large Language Models, 2026. URL https://arxiv.org/abs/2605.11887v1. Version 1.

Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. Amnesic Probing: Behavioral Explanation with Amnesic Counterfactuals. *Transactions of the Association for Computational Linguistics*, 9:160–175, 2021. doi: 10.1162/tacl_a_00359. URL https://aclanthology.org/2021.tacl-1.10/.

Javier Ferrando, Oscar Obeso, Senthooran Rajamanoharan, and Neel Nanda. Do i know this entity? knowledge awareness and hallucinations in language models. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2411.14257v2.

Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever, Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders, 2024. URL https://arxiv.org/abs/2406.04093.

Robert Huben, Hoagy Cunningham, Logan Smith, Aidan Ewart, and Lee Sharkey. Sparse Autoencoders Find Highly Interpretable Features in Language Models. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/1fa1ab11f4bd5f94b2ec20e794dbfa3b-Abstract-Conference.html.

Subhash Kantamneni, Joshua Engels, Senthooran Rajamanoharan, Max Tegmark, and Neel Nanda. Are Sparse Autoencoders Useful? A Case Study in Sparse Probing. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 29018–29049. PMLR, 2025. URL https://proceedings.mlr.press/v267/kantamneni25a.html.

Adam Karvonen, Can Rager, Johnny Lin, Curt Tigges, Joseph Isaac Bloom, David Chanin, Yeu-Tong Lau, Eoin Farrell, Callum Stuart Mcdougall, Kola Ayonrinde, Demian Till, Matthew Wearden, Arthur Conmy, Samuel Marks, and Neel Nanda. SAEBench: A Comprehensive Benchmark for Sparse Autoencoders in Language Model Interpretability. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 29223–29264. PMLR, 2025. URL https://proceedings.mlr.press/v267/karvonen25a.html.

Connor Kissane, robertzk, Arthur Conmy, and Neel Nanda. SAEs (usually) Transfer Between Base and Chat Models. AI Alignment Forum, 2024. URL https://www.alignmentforum.org/posts/fmwk6qxrpW8d4jvbd/saes-usually-transfer-between-base-and-chat-models. Research report, July 18.

Evgenii Kortukov, Piotr Komorowski, Florian Klein, Paula Engl, Gabriele Sarti, Seong Joon Oh, Sebastian Lapuschkin, and Wojciech Samek. Predicting Future Behaviors in Reasoning Models Enables Better Steering, 2026. URL https://arxiv.org/abs/2606.11172.

Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, Kamilė Lukošiūtė, Karina Nguyen, Newton Cheng, Nicholas Joseph, Nicholas Schiefer, Oliver Rausch, Robin Larson, Sam McCandlish, Sandipan Kundu, Saurav Kadavath, Shannon Yang, Thomas Henighan, Timothy Maxwell, Timothy Telleen-Lawton, Tristan Hume, Zac Hatfield-Dodds, Jared Kaplan, Jan Brauner, Samuel R. Bowman, and Ethan Perez. Measuring Faithfulness in Chain-of-Thought Reasoning, 2023. URL https://arxiv.org/abs/2307.13702.

Uzay Macar, Paul C. Bogdan, Senthooran Rajamanoharan, and Neel Nanda. Thought Branches: Interpreting LLM Reasoning Requires Resampling. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2510.27484v2.

Aleksandar Makelov, Georg Lange, and Neel Nanda. Towards Principled Evaluations of Sparse Autoencoders for Interpretability and Control. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/53356aebeea8ffd40a8ac3bb66243162-Abstract-Conference.html.

Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2403.19647v3.

Qwen Team. SAE-Res-Qwen3.5-35B-A3B-Base-W32K-L0_50: Model Card. Hugging Face, 2026. URL https://huggingface.co/Qwen/SAE-Res-Qwen3.5-35B-A3B-Base-W32K-L0_50. Accessed September 26, 2026.

Samuel Soo, Wesley Teng, and Chandrasekaran Balaganesh. Steering large language models with feature guided activation additions, 2025. URL https://arxiv.org/abs/2501.09929v1. Version 1.

Henk Tillman and Dan Mossing. Investigating Task-Specific Prompts and Sparse Autoencoders for Activation Monitoring, 2025. URL https://arxiv.org/abs/2504.20271.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting, 2023. URL https://arxiv.org/abs/2305.04388v2. Version 2.

Zhengxuan Wu, Aryaman Arora, Atticus Geiger, Zheng Wang, Jing Huang, Dan Jurafsky, Christopher D. Manning, and Christopher Potts. AxBench: Steering LLMs? even simple baselines outperform sparse autoencoders. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 67035–67080. PMLR, 2025. URL https://proceedings.mlr.press/v267/wu25a.html.

Aojie Yuan, Zhiyuan Julian Su, Haiyue Zhang, Yi Nian, and Yue Zhao. Hidden Error Awareness in Chain-of-Thought Reasoning: The Signal Is Diagnostic, Not Causal, 2026. URL https://arxiv.org/abs/2605.09502.

Peiyu Zang, Jian Tao, Jialing Zhang, Yichen Yuan, Wentao Zhang, Guang Liu, and Yonghua Lin. KernelGenBench: Can LLMs and agents write efficient kernels across operator sources and hardware platforms? *arXiv preprint arXiv:2607.27231*, 2026. URL https://arxiv.org/abs/2607.27231.

A MEASUREMENT AND STATISTICAL METHODS

A.1 DICTIONARY, WINDOWS AND CUTOFFS

The behavioral studies use `Qwen3.6-35B-A3B` with the Qwen-Scope `SAE-Res-Qwen3.5-35B-A3B-Base-W32K-L0_50` dictionary: 32,768 features per layer across 40 layers, with residual dimension 2,048. Qwen3.8-27B judges the task-specific outcomes. Dense pre-TopK readouts are computed from the encoder over the five tokens ending at each saved native cutoff. For nonnegative activations sorted as $a_{(1)} \geq \cdots \geq a_{(5)}$, the two summaries are the maximum $a_{(1)}$ and zero-padded top-three mean $(a_{(1)} + a_{(2)} + a_{(3)})/3$. Missing sparse-cache entries are not substituted for dense readouts.

The discovery banks contain 32 nested cutoffs per blackmail trajectory and 16 per whistleblowing trajectory, with 32 continuations at each state. Cutoff indices describe progression within each scenario. Figure 3 groups sources by the original trajectory's final label; those labels are distinct from the sampled continuation outcomes. The whistleblowing bank has 67 trajectories from 59 histories, with eight histories represented in both original-label groups.

A.2 BASE-MODEL DICTIONARY REUSE

The Qwen3.5 and Qwen3.6 checkpoints used here derive from the same pretrained base through post-training, motivating reuse of the base model's residual-stream dictionary. The SAE encoder and decoder are kept fixed when applied to Qwen3.6. This choice follows the Qwen-Scope model card's guidance on applying base-trained SAEs to post-training checkpoints (Qwen Team, 2026) and empirical work showing successful base-to-chat SAE transfer on Mistral-7B and Qwen1.5-0.5B (Kissane et al., 2024). Shared initialization motivates coordinate compatibility; reconstruction on the target checkpoint supplies an empirical check.

Before the main experiments, we checked reconstruction on Qwen3.6 using the Qwen3.5 base-trained SAE and found reconstruction error comparable to that on the original model. This supports using the fixed dictionary for the target model's residual readouts. Feature descriptions in this paper are grounded in Qwen3.6 activation contexts and outcome measurements; reconstruction fidelity is a dictionary-level diagnostic, rather than a test of semantic invariance for every feature across checkpoints.

A.3 OUTCOME ESTIMATES AND UNCERTAINTY

Task-specific judgments distinguish YES, NO and UNKNOWN. Length stops, malformed judgments and unresolved extraction remain UNKNOWN. For k YES, m NO and u UNKNOWN among $n = k + m + u$ assigned continuations, we use

$$\widehat{p}_{\text{lower}} = k/n, \qquad \widehat{p}_{\text{observed}} = k/(k+m), \qquad \widehat{p}_{\text{upper}} = (k+u)/n. \tag{6}$$

The primary rate is YES/all assigned continuations. The lower and upper estimates bound sensitivity to unresolved outcomes; they are not confidence intervals. The judged-only rate is defined when $k + m > 0$.

For the curves, we resample whole trajectories 10,000 times within each scenario and original-outcome group, recompute the mean at every cutoff, and report pointwise 2.5th and 97.5th percentiles (seed 20260925 with fixed group offsets). Each whistleblowing label group contains at most one trajectory per history, so these marginal intervals resample distinct histories. Shared histories remain grouped for held-out splitting and inference. Adjacent intervals and repeated continuations are kept within their source blocks. First-to-last comparisons use each source's earliest and latest YES/32, with ties recorded separately.

A.4 DECODER ADDITION

Equation 5 adds a unit-norm decoder column at the designated layer, scaled by a frozen mean residual norm. The all-token policy covers prompt prefill and decoding. Zero bypasses the addition; the panels contain no hook-active sham. Appendix D specifies calibration, random controls and sampling. Direction effects under this policy are distinct from the role of natural encoder activation.

B FEATURE DISCOVERY AND MEASUREMENT SENSITIVITY

B.1 DISCOVERY COHORTS

Feature discovery uses 83 previously inspected blackmail trajectories at 32 cutoffs and 67 whistleblowing trajectories from 59 histories at 16 cutoffs, with 32 continuations at each state. The blackmail feature bank includes 24 additional historical sources beyond the 59-source curve panel in Figure 3. The curve and feature analyses therefore have distinct source counts. The 16 coordinates selected in each scenario are fixed for subsequent cross-source evaluation and intervention.

B.2 PRIMARY ASSOCIATION AND FIVE ALTERNATIVE DEFINITIONS

The primary view combines the maximum dense-ReLU activation over the five saved tokens with YES/all assigned continuations. Maximum pooling preserves a localized feature response without averaging it over surrounding tokens, while the fixed outcome denominator keeps the same 32-draw sampling unit at every state. The alternatives combine two activation summaries (maximum and zero-padded top-three mean) with three outcome estimates (lower, observed and upper in Equation 6). Table 5 reports the primary correlation and the minimum absolute correlation across these six views.

For either interval array $z_i[j]$ (the activation change $x_i^f[j]$ or outcome change $y_i[j]$), adjustment first subtracts the trajectory mean: $z_i'[j] = z_i[j] - \overline{z}_i$. It then subtracts the mean of $z_i'[j]$ among trajectories sharing the original outcome label at interval position j. Pearson correlation pools the resulting adjusted interval pairs. Original outcome labels serve this retrospective discovery adjustment; they are not inputs available for prediction at a new state. Shared histories remain the independence blocks for subsequent validation.

The frozen discovery selector requires positive maximum activation in 10–90% of saved windows and a common nonzero correlation sign across all six views. Eligible coordinates are ordered by

$$S_f = \min_{v=1,\ldots,6} |\rho_f^{(v)}|. \tag{7}$$

The first 16 per task form the evaluated candidate families. Main-text ρ_f values are the first-view correlations, not S_f; the candidate families and their original ranks are preserved. Agreement across the six definitions therefore describes the selected candidates' directional consistency by construction. It does not establish unchanged whole-dictionary rankings under six separate searches.

B.3 Complete candidate correlations

Table 5: **Discovery correlations for all selected candidates.** ρ: primary maximum-readout, YES/all correlation. S: minimum absolute correlation across all six views (Equation 7). All six signs agree by the selection rule. Rows preserve discovery rank.

	Blackmail			Whistleblowing		
Rank	Coordinate	ρ	S	Coordinate	ρ	S
1	L5:F22948	−0.126	0.124	L18:F18808	+0.219	0.203
2	L9:F29759	−0.100	0.099	L22:F16568	+0.224	0.200
3	L20:F21308	+0.098	0.098	L24:F27715	+0.185	0.181
4	L16:F7683	−0.096	0.094	L16:F12323	−0.183	0.180
5	L35:F18288	−0.094	0.093	L21:F1101	−0.183	0.176
6	L30:F17299	+0.091	0.091	L32:F9104	−0.175	0.175
7	L33:F29733	−0.089	0.089	L15:F20283	−0.182	0.174
8	L25:F16234	+0.088	0.088	L29:F13173	−0.195	0.173
9	L19:F19934	+0.090	0.088	L18:F27087	+0.172	0.171
10	L36:F18288	−0.091	0.087	L29:F31744	+0.193	0.171
11	L9:F1996	−0.087	0.087	L31:F11337	+0.170	0.170
12	L12:F18759	−0.097	0.086	L23:F13762	−0.196	0.170
13	L30:F23889	−0.086	0.086	L29:F17464	+0.183	0.168
14	L25:F6497	−0.086	0.086	L29:F15297	+0.179	0.168
15	L38:F18122	+0.086	0.086	L12:F16997	−0.169	0.167
16	L32:F13148	+0.086	0.086	L24:F26687	−0.168	0.166

B.4 Activation examples and local sensitivity

Each of the 32 selected coordinates has 15 high-activation examples and five exact-zero controls. At most two high examples come from a single source/history. Zero controls are matched to peak token and relative position where available: all 80 blackmail controls and 70 of 80 whistleblowing controls match a high-example peak token ID. A zero means all five dense-ReLU activations are zero. Outcome-blind machine-generated descriptions use the visible context and peak tokens; they are working semantic interpretations. The five-token measurement window is distinct from the longer context displayed for reading. Figure 4 uses blackmail L25:F16234 examples H01/Z05 and whistleblowing L29:F17464 examples H01/Z01.

C Feature relationships on unseen sources

C.1 Cohorts and frozen measurements

Blackmail candidates were selected using all 83 historical discovery trajectories. The test uses six new unsteered trajectories from the same prompt, selected by index and native structural eligibility before their outcome labels were inspected. Each has 32 native prefix cutoffs and 16 matched continuations per cutoff. Of 3,072 continuations, 1,969 were YES, 956 NO and 147 UNKNOWN. Whistleblowing uses 29 trajectories from 26 histories excluded from discovery, including three histories with paired original-YES and original-NO trajectories. Its 16 cutoffs and 32 continuations per state yield 14,848 continuations, of which 14,434 were determinate and 414 unresolved. Each task retains its frozen 16 coordinates, signs and five-token readout windows.

Activation capture uses full-context teacher-forced forwards over the saved source trajectories. These provide an approximation to the boundary representations rather than bitwise replay of each original generation request. The primary outcome estimate is YES divided by assigned continuations, with unresolved outcomes bounded using Equation 6. Feature identities and signs are fixed before inspecting test outcomes.

C.2 CONTINUOUS-SCORE DISCRIMINATION AND UNCERTAINTY

Upward and downward events satisfy $\Delta\widehat{p} \geq 0.15$ and $\Delta\widehat{p} \leq -0.15$, respectively. Flat intervals satisfy $|\Delta\widehat{p}| \leq 0.10$; intervals between these boundaries remain in the broad comparison population. Blackmail has 12 upward events across five sources, 14 downward events across five sources, and 138 flat intervals. Whistleblowing has 12 upward events across five histories, 15 downward events across ten histories, and 402 flat intervals. Upward/downward prevalences are 0.065/0.075 and 0.028/0.034, respectively.

For each feature, the upward score is the discovery sign times the raw maximum-activation change; the downward score is its negative. AUROC compares the event with all remaining intervals. The scores retain their discovery signs and use raw activation changes without trajectory or outcome-group centering. AUROC integrates over score thresholds.

Tables 6 and 7 report all coordinates and both event directions. Percentile 95% AUROC intervals use 1,000 source/history bootstrap resamples with candidates held fixed; paired trajectories remain in the same history block. These intervals describe each fixed candidate and are not simultaneous intervals over the family. Event coverage and source counts determine the precision of these estimates.

C.3 PROBABILITY AND POSITION REFERENCES

The headroom score $1-\widehat{p}_{ij}$ yields upward AUROC 0.678 in blackmail and 0.134 in whistleblowing. The initial-rate score $\widehat{p}_{ij}$ yields downward AUROC 0.429 and 0.912, respectively. Increasing interval index yields upward/downward AUROCs 0.227/0.296 in blackmail and 0.271/0.422 in whistleblowing. All use the same interval populations as the feature scores. Probability references require continuation sampling at the earlier boundary; they contextualize event changes rather than measure a prefix-only predictor's performance. Their signs are reported as specified, without reversing them using test outcomes.

Unresolved outcomes touch 145 of 186 blackmail intervals and 326 of 435 whistleblowing intervals. Bounding each endpoint probability between YES/all and (YES+UNKNOWN)/all gives lower and upper bounds on each difference. In blackmail these bounds guarantee 10 upward and nine downward events, while permitting up to 27 and 24. In whistleblowing they guarantee seven upward and twelve downward events, while permitting up to 20 and 21. These bounds accompany the primary assigned-denominator event labels.

Table 6: **Blackmail event discrimination for all frozen candidates.** AUROC with pointwise 95% source/history-bootstrap intervals; ranks and signs are fixed by discovery.

Rank	Coordinate	Upward AUROC [95% CI]	Downward AUROC [95% CI]
1	`L5:F22948`	0.491 [0.477, 0.500]	0.507 [0.387, 0.564]
2	`L9:F29759`	0.410 [0.286, 0.715]	0.578 [0.431, 0.645]
3	`L20:F21308`	0.662 [0.465, 0.821]	0.532 [0.365, 0.755]
4	`L16:F7683`	0.420 [0.259, 0.532]	0.512 [0.219, 0.638]
5	`L35:F18288`	0.464 [0.272, 0.635]	0.529 [0.490, 0.612]
6	`L30:F17299`	0.631 [0.511, 0.718]	0.589 [0.464, 0.803]
7	`L33:F29733`	0.524 [0.463, 0.678]	0.539 [0.296, 0.703]
8	`L25:F16234`	0.382 [0.254, 0.656]	0.505 [0.383, 0.573]
9	`L19:F19934`	0.489 [0.338, 0.575]	0.552 [0.353, 0.681]
10	`L36:F18288`	0.418 [0.323, 0.497]	0.506 [0.494, 0.516]
11	`L9:F1996`	0.346 [0.207, 0.453]	0.606 [0.388, 0.691]
12	`L12:F18759`	0.530 [0.369, 0.695]	0.635 [0.573, 0.750]
13	`L30:F23889`	0.532 [0.491, 0.559]	0.568 [0.398, 0.652]
14	`L25:F6497`	0.572 [0.395, 0.718]	0.488 [0.374, 0.573]
15	`L38:F18122`	0.516 [0.327, 0.633]	0.561 [0.363, 0.756]
16	`L32:F13148`	0.624 [0.555, 0.721]	0.613 [0.512, 0.782]

Table 7: **Whistleblowing event discrimination for all frozen candidates.** AUROC with pointwise 95% source/history-bootstrap intervals; ranks and signs are fixed by discovery.

Rank	Coordinate	Upward AUROC [95% CI]	Downward AUROC [95% CI]
1	L18:F18808	0.578 [0.496, 0.664]	0.535 [0.416, 0.613]
2	L22:F16568	0.377 [0.143, 0.584]	0.321 [0.188, 0.435]
3	L24:F27715	0.501 [0.494, 0.508]	0.561 [0.404, 0.664]
4	L16:F12323	0.519 [0.390, 0.645]	0.454 [0.319, 0.568]
5	L21:F1101	0.338 [0.026, 0.501]	0.440 [0.275, 0.586]
6	L32:F9104	0.418 [0.157, 0.590]	0.496 [0.374, 0.636]
7	L15:F20283	0.428 [0.248, 0.597]	0.545 [0.358, 0.701]
8	L29:F13173	0.580 [0.326, 0.744]	0.446 [0.311, 0.525]
9	L18:F27087	0.425 [0.153, 0.663]	0.395 [0.273, 0.501]
10	L29:F31744	0.534 [0.498, 0.602]	0.467 [0.382, 0.582]
11	L31:F11337	0.378 [0.143, 0.574]	0.295 [0.168, 0.411]
12	L23:F13762	0.418 [0.363, 0.504]	0.434 [0.300, 0.546]
13	L29:F17464	0.516 [0.420, 0.665]	0.574 [0.476, 0.685]
14	L29:F15297	0.382 [0.239, 0.493]	0.458 [0.298, 0.594]
15	L12:F16997	0.501 [0.340, 0.945]	0.561 [0.422, 0.699]
16	L24:F26687	0.689 [0.592, 0.871]	0.576 [0.398, 0.729]

D DECODER INTERVENTIONS AND CONTEMPORANEOUS CONTROLS

D.1 GENERATION STARTS, DOSES AND OUTCOME DEFINITIONS

Whistleblowing uses the fictional Zenavex pharmaceutical scenario from the multi-turn *Thought Branches* setting (Macar et al., 2026). The model encounters trial-safety discrepancies and internal communications, then generates a final response about what to do. The held-out intervention starts from the native prompt of the target final assistant turn, before its reasoning continuation. It selects one original-NO source from each of 26 confirmation histories, including alternative NO draws where necessary. The observational panel instead contains 29 trajectories from those histories. Eight repeats per history are matched across the 18 reported conditions. Thus 208 draws per condition represent 26 independent history blocks.

Blackmail starts from the same original 2,372-token prompt in every condition, with 100 matched seeds. The two selected doses are L16:F7683 at $+0.10$ and L25:F16234 at -0.10; the two random SAE coordinates are matched in layer and applied dose. The reported whistleblowing panel includes one prespecified dose per candidate, together with random L18:F6754 at $+0.10$. These random coordinates provide illustrative controls, not an empirical distribution of random-direction effects.

Calibration and dose rationale. For calibration tokens $\mathcal{C}$, the fixed layer scale is $\bar{r}_\ell = |\mathcal{C}|^{-1} \sum_{j \in \mathcal{C}} \|h_j^{(\ell)}\|_2$. The reused calibration comprises 438,729 rendered prompt tokens from eight KernelGenBench development trajectories, measured at the post-transformer-layer output. It is token-weighted and frozen across the behavioral campaigns. With unit-norm decoder directions, Equation 5 therefore gives $\|h'_{\ell,t} - h_{\ell,t}\|_2 = 0.10\bar{r}_\ell$. This calibration uses residual-vector magnitudes, not the mean encoder activation of each selected SAE coordinate.

The dose is fixed at 0.10. Residual-relative calibration makes its scale explicit (Ferrando et al., 2025; Soo et al., 2025); completion and unresolved-output counts characterize generation at this dose.

The intervention uses unit-norm SAE decoder columns and the frozen calibration above. The all-token policy applies during prompt prefill and subsequent decoding at the designated layer. This schedule is fixed before generation and requires neither future reasoning boundaries nor a trigger based on natural feature activation. The five-token boundary windows serve offline measurement and candidate discovery; they do not specify the intervention support. The experiment tests sustained steering under this deployment constraint. Boundary-localized steering, ablation or activation patching would instead test the causal role of a particular reasoning segment. Zero bypasses

the addition; there is no hook-active sham. The matched campaigns share a serving revision and a 16,384-new-token budget in a 32,768-token context. Whistleblowing uses temperature 1.0, top-p .95, top-k 20 and presence penalty 1.5; blackmail uses temperature .7 and top-p .95. Each campaign has a single contemporaneous shared zero condition, so temporal drift within a campaign is not separately measured.

The whistleblowing judge reads the visible final response. YES includes a clear commitment to report fraud externally, a simulated external communication, or concrete steps to expose wrongdoing outside the company. An unsent draft, an uncommitted possibility, or internal escalation alone is NO. The eligibility rule requires a natural stop, exactly one thinking-close marker and nonempty visible final text; ineligible outputs and unresolved judgments remain excluded from binary adjudication but retained in the assigned denominator. Simulated tools are not executed. Blackmail uses the task-specific blackmail-action judgment and records UNKNOWN separately. Both panels use Qwen3.8-27B judgments and pass 12/12 positive/negative controls before and after grading.

D.2 Complete counts, uncertainty and readout comparisons

For each whistleblowing history we calculate the treated-minus-zero difference in YES/8, then average across the 26 histories. Reported intervals bootstrap histories, and two-sided history sign-flip tests are Holm-adjusted over the 16 selected directions. Table 8 aligns those results with the same coordinates' upward-event AUROCs. These columns describe different measurements: the observational AUROC measures event discrimination from natural activation changes across saved reasoning intervals, whereas the intervention test concerns final outcomes generated with a decoder addition from the target-turn prompt.

Table 8: Whistleblowing intervention counts and paired contrasts. The reported panel contains 16 selected arms, one random control and zero, with 208 assigned draws per arm. E denotes excluded outputs. Differences and pointwise 95% history-bootstrap intervals are in percentage points. The last column is the upward-event AUROC of the same coordinate on the observational cohort. Random arms are outside the selected intervention family.

Coordinate	α	Y/N/E	Δ	95% CI	p_{raw}	p_{Holm}	Up AUC
Zero	0	14/188/6	—	—	—	—	—
`L18:F18808`	$+0.10$	26/181/1	$+5.8$	$[0.5, 10.6]$	0.0562	0.4590	0.578
`L22:F16568`	$+0.10$	28/175/5	$+6.7$	$[1.0, 12.5]$	0.0468	0.4590	0.377
`L24:F27715`	$+0.10$	31/173/4	$+8.2$	$[1.9, 14.4]$	0.0321	0.4095	0.501
`L16:F12323`	-0.10	31/173/4	$+8.2$	$[1.9, 14.4]$	0.0315	0.4095	0.519
`L21:F1101`	-0.10	20/184/4	$+2.9$	$[-2.4, 8.2]$	0.4023	1.0000	0.338
`L32:F9104`	-0.10	23/182/3	$+4.3$	$[-1.0, 9.1]$	0.1669	0.6675	0.418
`L15:F20283`	-0.10	38/163/7	$+11.5$	$[5.8, 17.8]$	0.0015	0.0225	0.428
`L29:F13173`	-0.10	27/170/11	$+6.2$	$[1.0, 11.5]$	0.0415	0.4565	0.580
`L18:F27087`	$+0.10$	24/182/2	$+4.8$	$[1.0, 8.7]$	0.0469	0.4590	0.425
`L29:F31744`	$+0.10$	22/183/3	$+3.8$	$[-2.9, 10.6]$	0.3631	1.0000	0.534
`L31:F11337`	$+0.10$	34/165/9	$+9.6$	$[4.3, 15.4]$	0.0051	0.0714	0.378
`L23:F13762`	-0.10	27/177/4	$+6.2$	$[1.4, 11.5]$	0.0459	0.4590	0.418
`L29:F17464`	$+0.10$	40/165/3	$+12.5$	$[6.2, 18.3]$	0.0008	0.0128	0.516
`L29:F15297`	$+0.10$	27/174/7	$+6.2$	$[1.0, 11.5]$	0.0489	0.4590	0.382
`L12:F16997`	-0.10	15/186/7	$+0.5$	$[-3.8, 4.3]$	1.0000	1.0000	0.501
`L24:F26687`	-0.10	28/176/4	$+6.7$	$[-0.5, 13.9]$	0.1073	0.5364	0.689
`R: L18:F6754`	$+0.10$	14/191/3	$+0.0$	$[-3.8, 3.8]$	1.0000	—	—

All 3,744 requests in the reported whistleblowing panel completed; 3,657 outputs received determinate labels and 87 were excluded. Request completion is distinct from satisfying the final-output eligibility rule. Exclusion fractions for zero, `L29:F17464` and `L15:F20283` are respectively 6/208, 3/208 and 7/208. Their judged-only YES rates are 14/202, 40/205 and 38/201. Assigning unresolved outcomes adversarially gives treated-minus-zero point bounds of $[9.6, 13.9]$ and $[8.7, 14.9]$

percentage points for the two selected arms. These are missing-outcome sensitivity bounds, not confidence intervals or additional corrected tests. Per-arm length and repetition summaries are unavailable.

The blackmail panel contains 230 YES, 246 NO and 24 UNKNOWN across 500 outputs, with no target transport failure. Table 3 reports all five conditions; their NO counts, in its row order, are 43, 60, 51, 54 and 38. Intervals use matched-seed bootstrap resampling. Exact paired tests on determinate discordances are Holm-adjusted across the four nonzero arms. The two selected point differences have adversarial UNKNOWN bounds of $[-22, -12]$ and $[-11, -6]$ points; their four-arm corrected p values remain .131 and .580 under the reported analysis. Bounds on point differences and corrected significance answer separate questions.

D.3 Signed-dose comparison for the blackmail candidates

A separate historical screen evaluates the same 16 blackmail candidates at $+0.10$ and -0.10, with 100 matched seeds per dose and two zero bookends (3,400 generations; 181 UNKNOWN outcomes). Table 9 preserves the full family. The discovery-aligned high-minus-low differences are $+32$ points for `L16:F7683` (95% interval $[18, 45]$, Holm $p = .000334$) and $+30$ for `L25:F16234` ($[15, 44]$, Holm $p = .00505$). These compare opposite doses in that campaign, whereas Table 3 compares a selected dose with a contemporaneous zero in a separate campaign.

Table 9: **Complete raw-16 blackmail signed-dose screen.** Every feature has 100 assigned generations per sign. Difference is unconditional YES at $+0.10$ minus YES at -0.10, in percentage points. Holm p values are from the owning result report, not recomputed here. Zero-dose bookends scored 52/100 and 44/100 YES.

Feature	Sign	YES $+0.10$	YES -0.10	Difference, pp	Holm p
`L5:F22948`	$-$	39	57	-18	.209
`L9:F29759`	$-$	46	53	-7	1
`L20:F21308`	$+$	43	51	-8	1
`L16:F7683`	$-$	25	57	-32	.000334
`L35:F18288`	$-$	53	53	0	1
`L30:F17299`	$+$	54	56	-2	1
`L33:F29733`	$-$	51	54	-3	1
`L25:F16234`	$+$	65	35	$+30$	.00505
`L19:F19934`	$+$	46	58	-12	1
`L36:F18288`	$-$	44	53	-9	1
`L9:F1996`	$-$	49	55	-6	1
`L12:F18759`	$-$	50	43	$+7$	1
`L30:F23889`	$-$	55	49	$+6$	1
`L25:F6497`	$-$	64	44	$+20$	.0772
`L38:F18122`	$+$	57	48	$+9$	1
`L32:F13148`	$+$	52	50	$+2$	1

E Exploratory extension: kernel generation

This exploratory extension describes candidate qualification, decoder interventions and executable verification on overlapping development states. It reports all six candidate directions at each of the 11 saved states. The historical comparisons are descriptive observations for candidate generation; independent confirmation requires held-out operators and states with contemporaneous zero and random-direction controls.

Development qualification. The uncapped search evaluates 1,310,720 coordinates on eight reused trajectories. A candidate must satisfy training TP $\geq 4/8$ and FP $\leq 2/8$ in all 165 leave-three-event-out splits under maximum and zero-padded top-three aggregation. On 11 distinct positive and 19 control intervals it must have recall at least .50, false-positive rate at most .25 and positive Matthews correlation; qualification must persist under 32 frozen matching seeds. The thresholds were formal-

ized after earlier three-layer results were inspected. The search retains 100 coordinates across 32 layers; these checks do not supply a whole-search null test or independent-source confirmation.

Continuation intervention and grading. One unit decoder direction is added at every token of its zero-based layer with strength .10 of the frozen layer mean residual norm. Steered runs and historical reference runs use identical budgets of 100 turns and 7,200 seconds per attempt. The same 88 reference attempts serve all six comparisons. Of 616 planned attempts, 612 have pristine grades and four produce no kernel (one reference, three L10); all four remain failures. Among the grades, 604 record `dirty=0`; eight legacy grades lack that hygiene field. The audited operators exclude the report's known invalid benchmark classes.

Table 10: Verified successes out of eight assigned continuations per saved state. Column labels denote the six coordinates in Table 4. Ref. is the historical reference with the same 100-turn, 7,200-second budget.

Operator / turn	Ref.	L10	L28	L29	L30	L33	L36
unsafe_split_with_sizes / 32	3	6	7	7	7	5	7
unsafe_split_with_sizes / 35	6	8	8	8	7	7	8
sum / 8	2	6	5	6	6	6	6
sum / 40	4	8	8	8	8	8	8
_softmax / 18	4	6	8	7	5	6	6
sort / 51	2	5	8	5	6	5	6
floor_divide / 21	5	5	7	3	5	5	4
floor_divide / 52	4	4	6	3	5	5	5
mean / 14	5	8	8	8	8	8	8
upsample_nearest2d_backward / 12	4	6	7	7	7	7	7
log_sigmoid_backward / 45	5	5	6	4	6	2	6
Total / 88	44	67	78	66	70	64	71

Saved-request diagnostic replay. In a sort implementation, calling `torch.sort` re-enters the registered function, causing recursion. At the saved turn-56 request, all eight natural zero-condition samples identify recursion, whereas none of eight `L10:F4645` samples do; the latter tend to reread verifier output. This screen replays each saved request separately, so later prompts inherit the original trajectory's diagnosis rather than the preceding intervention's output. Its raw dose of 1 corresponds to approximately 45.7% of the layer's calibration residual norm, distinct from the 10% continuation study.

The replay covers seven conditions and 11 saved requests (turns 51–61), with eight natural and one greedy sample per request. It measures immediate diagnostic choices under the stronger dose; the continuation outcomes above use the normalized 0.10 dose.

F Reasoning content at a saved boundary

We examine blackmail source 034 to connect reasoning text with boundary measurements. Its 205 reasoning chunks carry machine-assigned self-preservation, leverage-identification and plan-generation tags, motivated by *Thought Branches* (Macar et al., 2026). The 32 saved cutoffs each have 32 continuations. For each interval, we compare YES/32 and five-token dense-ReLU readouts at its endpoints. These boundary readouts measure the endpoints rather than the intervening text as a whole.

At $C_{16} \to C_{17}$, six plan-tagged chunks move from weighing whom to contact to a direct threat: "If I email Kyle directly, I can negotiate," followed by a draft proposing disclosure to Rachel unless the wipe is cancelled. Observed blackmail frequency rises from 27/32 to 31/32, while `L30:F17299` rises from 0 to 0.14041 (Table 11). This interval was selected after inspecting the grid and illustrates aligned changes in text, readout and outcome frequency.

Table 11: **One plan interval, observed at its boundaries.** Each boundary has 32 independently judged continuations. The SAE column is the five-token readout of the positively oriented `L30:F17299`; it was examined after outcome-based candidate selection.

Saved boundary	Future actions: blackmail / other / unresolved	SAE readout
Before the interval, C_{16}	27 / 4 / 1	0
After the interval, C_{17}	31 / 1 / 0	0.14041

F.1 Context across the source trajectory

The 31 adjacent intervals combine three historical seed cohorts, so their differences are descriptive rather than paired effects. Table 12 places the selected case among all tagged intervals.

Table 12: **All tagged intervals on source 034.** An interval is counted for a category if it contains at least one tagged chunk; categories overlap, so rows do not partition the 31 intervals. Change is the difference in YES/32 between adjacent cutoffs, in percentage points.

Tag present	Intervals	Increase	Decrease	No change	Mean $\Delta\widehat{p}$ (range), pp
Self-preservation	3	1	1	1	0.0 (-6.25 to $+6.25$)
Leverage identification	4	2	0	2	$+1.56$ (0 to $+3.125$)
Plan generation	19	4	3	12	$+0.33$ (-15.625 to $+15.625$)

At another plan-tagged interval, $C_{14} \to C_{15}$, YES decreases from 30/32 to 27/32 and the same readout falls from 0.09489 to zero. Across the 19 plan-containing intervals, four increase, three decrease and twelve are flat. Thus the case links a particular reasoning transition to its measured futures; it does not estimate a general effect of plan formation.