COGNITIVE STEERING OF LANGUAGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Current approaches to identifying representations that influence a language model's (LM) behavior often rely on noisy datasets or imprecise linguistic descriptions to define complex behavioral profiles. Such approaches can be confounded by spurious features of the data and our own biases in defining their content, resulting in unexpected behaviors at intervention time. In cognitive science, probabilistic generative models of cognition provide formal accounts of variables that parameterize complex human behaviors, including behaviors relevant to LM personality and alignment. In this work, we treat cognitive models as data-generating processes and use their outputs to supervise the search for representations corresponding to their parameters, which we term *cognitive variables*. We identify these representations in both activation and weight space and use them for *cognitive steering*: controlling LM behavior according to the computations assumed by the originating cognitive models. According to LM judges, cognitive variables derived from models of interpersonal communication control free generations in naturalistic settings beyond the toy domains the cognitive models were designed to represent, most reliably when distilled into model weights. Our approach provides a principled method for steering naturalistic LM behavior using empirically validated, formal models of human cognition.

1 INTRODUCTION

Mechanistic interpretability is motivated by the prospect of control: if we understood how a model's internal state gives rise to its behavior, we could intervene on that state to produce the behavior we care about (Mueller et al., 2025). Recent techniques allow us to identify circuits Conmy et al. (2023) and steer models (Rimsky et al., 2024; Arditi et al., 2024) with remarkable precision. But our ability to say exactly *what* we want them to do has lagged behind, especially when the desired behaviors are high-level properties we notice in people, such as being moral, considerate, or polite. In other words, *the bottleneck is no longer the precision of our interventions but the precision of our behavioral expectations.*

Target behaviors are typically specified through natural-language descriptions of desired features in the system prompt, or through sets of contrastive demonstrations which are themselves often generated from such descriptions (Chen et al., 2025; Lu et al., 2026; Li et al., 2023; Rimsky et al., 2024). These methods are effective for many purposes, but begin to show strain when the target behavior grows more abstract. Demonstrations may vary along more dimensions than the one we intend, and descriptions carry our own unstated assumptions. Consequently, interventions learned through a contrastive regime are empirically underdetermined, absorbing spurious features of the examples and changing the LM in ways we did not intend (Tan et al., 2024).

Probabilistic generative models designed to explain human behavior, known as *cognitive models*, provide candidate formalisms for specifying high-level behaviors more precisely. Cognitive models are designed to explain how high-level psychological traits arise from interactions among interpretable variables: the costs, goals, and rewards an agent weighs in choosing how to act. Such models have formalized behaviors ranging from pragmatic communication to social reasoning; a model of polite speech (Yoon et al., 2020), for instance, describes how a speaker balances informational and social goals, to produce honest, kind, or sycophantic utterances depending on their relative weights. Prior work fitting such models to LM behavior has found that their parameters respond systematically to manipulations of system prompts, reasoning budgets, alignment datasets, and base models (Murthy et al., 2025; 2026). These results provide a rigorous account of socially

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.

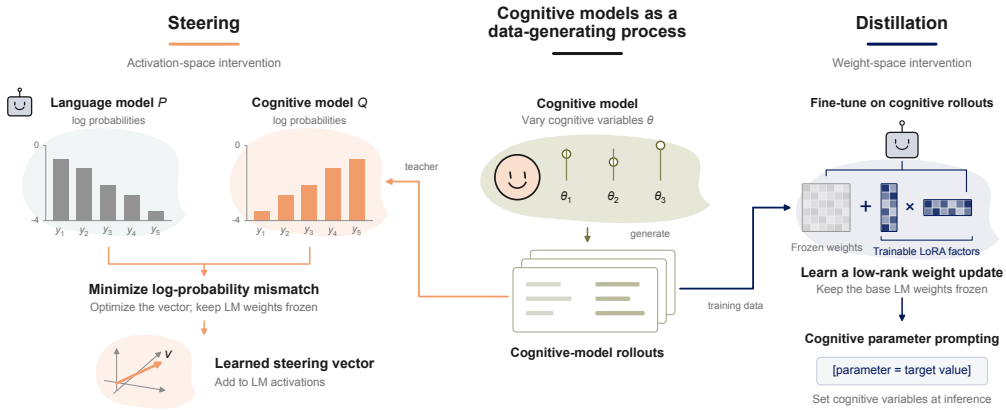

Figure 1: **Overview of our cognitive steering approach.** We operationalize cognitive models as data generating sources of supervision for language models, providing the training data for both (left) steering vectors that are used for test-time behavioral interventions in activation space and (right) distillation fine-tuning for train-time behavioral interventions. We select a set of cognitive models that provide dense coverage over the phenomenon of interpersonal communication, and use their parameters that capture a variety of tradeoffs between social goals as targets for steering.

relevant LM behavior at the output level, and raise a natural next question: how are the corresponding variables represented *inside* the LM?

Work connecting computational accounts of behavior to LM representations (cf. Ku et al., 2025) has proceeded from two relevant perspectives. Methods for identifying graph-like accounts of behavior in LM representations rely on causal graphs that typically describe toy concepts rather than the complexity captured by human-validated models of behavior (Geiger et al., 2024; Conmy et al., 2023). Conversely, work using cognitive models for output-level interpretability, including a growing line that uses probabilistic models to account for LM control (Bigelow et al., 2026), has generally not connected their computational variables to representations inside the LM (Bigelow et al., 2024; Schubert et al., 2024). We bridge these perspectives by treating cognitive models as *data-generating processes*. Sampling from the full range of a cognitive model's parameters guides the search for *cognitive variables*, representations in the LM's activation and weight space whose manipulation produces the graded behavioral changes that the parameters produce in the cognitive model.

Contributions We consider a range of activation-space and weight-space interventions guided by rollouts from probabilistic generative models of human cognition. In weight space, we fine-tune a base LM on rollouts from a cognitive model fit to human data and recover cognitive variables at inference time through simple prompting. In activation space, we use the same data to learn steering vectors and affine maps for individual cognitive-model parameters. **Within the cognitive models' own domains**, optimized vectors, affine maps, and distillation all recover target parameters far more accurately than the unsteered LM or prompting. **We then evaluate these methods in naturalistic, out-of-distribution scenarios of interpersonal communication**, whose dynamics involve the trade-offs specified by cognitive models of pragmatic reasoning and social cognition. There, **distillation shifts LM-judged qualities of free generations in the target direction** for 13 of 16 parameters in both `Llama-3.1-8B-Instruct` and `Qwen3-8B`, while activation steering transfers less consistently. In a case study on self-disclosure in a dating profile scenario, we also observe that steering changes concrete choices. Finally, model parameters without a natural counterpart in an LLM, such as a prior inclination to stay silent, resist steering, suggesting that *cognitive steering* might be helpful in identifying which variables of human behavior an LLM shares.

2 Cognitive Models as a Data-Generating Process

Cognitive models constitute a large and heterogeneous family with different philosophical and computational commitments, ranging from richly structured world models that support inference, learn-

ing, and decision-making (Craik, 1943; Johnson-Laird, 1980; Gentner & Stevens, 1983; Ullman et al., 2017; Ullman & Tenenbaum, 2020) to simpler accounts that predict behavior via bounded search, satisficing, and adaptive heursitics (Simon, 1955; Newell & Simon, 1972; Payne et al., 1988; Gigerenzer & Goldstein, 1996); cue and feature matching (Brunswik, 1952; Medin & Schaffer, 1978; Tversky, 1977; Nosofsky, 1986; Dawes, 1979); or formal operations over symbolic representations (Rips, 1994).

We set these philosophical commitments aside, and adopt a few operational criteria for selecting cognitive models that are able to be used for cognitive steering. First, we use cognitive models that specify a conditional generative process over behavior. This allows us to sample observable behaviors conditional on a context (world state or stimulus), and is in contrast to cognitive models that merely return point predictions, rankings, or scores. Second, we use cognitive models whose latent variables retain stable interpretations across contexts. In other words, we use cognitive models whose distributions over behaviors can represent variation in relevant *latent* cognitive variables, such as an agent's goals or value function, rather than variation in *observable* features of the world or stimulus that condition those distributions. [1] Finally, among models that satisfy these properties, we intervene on the structural weights in the decision rule as specified by the original model, rather than quantities the paper fixes, derives, or computes from elicited data (e.g., rationality or temperature parameters).

2.1 Model selection

For this work, we use cognitive models of human interpersonal communication to investigate cognitive steering in LMs. To explore this large space, we conduct a LM-driven literature search through relevant venues and then perform a three-stage filtering process. First, the LM is prompted to identify published models that abide by our above-mentioned operational criteria. Second, we filter the candidate models on features that are unique to the interpersonal communication settings we study. These include (a) that the model's choices y are made over a space of utterances, rather than actions or scalar responses, and (b) that the parameter(s) we intervene on are weights in a speaker's own trade-off between particular goals. We assume that the LM is taking the position of a speaker in a speaker-listener dyad. As a final filter, a subset of the authors independently rated each model's thematic relevance to out-of-distribution interpersonal communication scenarios[2], yielding the final set of nine cognitive models in Table 1 (see Section A.1 for further implementation details).

Selected cognitive variables The resulting set varies in both structure and parameters (Table 1). Six of the nine models are rational speech act models of pragmatic reasoning, and the others model indirect communication, persona construction, and pedagogical advice. Their 16 parameters span politeness, perspective-taking, impression management, affect, and advice giving, and we select them for dense coverage over the fundamental processes of interpersonal communication (Berger & Roloff, 2019) that our out-of-distribution scenarios draw on (Section 4). The models also differ in how their parameters encode trade-offs between goals. In three models, a single parameter weighs two goals against each other (*within-parameter*; e.g., RSA-08's λ trades relevance against truthfulness). In four, each parameter weights one goal and competes with the model's other parameters (*across-parameters*): the weights of RSA-02, RSA-03 and RSA-07 sum to one, and LANG-02's persona values compete in the speaker's choice of persona. The remaining two parameters, a planning horizon for advice and a prior inclination to stay silent, are not traded against any other goal.

[1] In econometrics, this is understood as a distinction between "policy invariant" models whose *structural parameters* are meant to be invariant to the stimulus and thus support counterfactual intervention with the stimulus held fixed as opposed to "reduced-form equations" that describe statistical patterns of variance and covariance in data (Haavelmo, 1944; Marschak, 1953; Hurwicz, 1962; Lucas, 1976). In cognitive science, such structural parameters point to a mechanistic account of information processing as it is done by a mind.

[2] We note that this human validation process is likely under-powered and plan to address this step in future work.

Model	Parameter	Definition	Trade-off
IMPR-03 (Achimova et al., 2025)	γ	express your opinion rather than avoid conflict with the listener's opinion	within-parameter
RSA-08 (Sumers et al., 2024)	λ	be relevant rather than truthful	
RSA-09 (Mozuraitis et al., 2018; Ryskin et al., 2020)	$p_{\text{addressee}}$	describe things from the listener's point of view rather than your own	
LANG-02 (Burnett, 2019)	μ(stern leader)	be seen as competent but cold	across-parameters
	μ(doofus)	be seen as friendly but not competent	
	μ(cool guy)	be seen as competent and friendly	
RSA-02 (Yoon et al., 2020)	ω_{inf}	be accurate and informative	
	ω_{soc}	make the listener feel good	
	ω_{pres}	be seen as both honest and kind	
RSA-03 (Carcassi & Franke, 2023)	ω_{inf}	be accurate and informative	
	ω_{soc}	make the listener feel good	
RSA-07 (Kao & Goodman, 2015)	w_{state}	convey what the situation is actually like	
	w_{valence}	convey whether you feel good or bad	
	w_{arousal}	convey how strongly you feel	
PED-02 (Sumers et al., 2022)	H	number of future choices the advice should serve	no trade-off
RSA-11 (Qing et al., 2016)	$\pi_{\varnothing}$	prior inclination to say nothing	

Table 1: **Cognitive model suite and target parameters for steering.** The trade-off column describes how a model's parameters set the balance between competing goals. *Within-parameter*: one parameter weighs two goals against each other, so raising it favors one goal at the expense of the other. *Across-parameters*: each parameter weights a single goal, and raising it shifts weight away from the model's other goals. *No trade-off*: the parameter scales a single behavior and is not weighed against another goal.

3 METHODS

3.1 PROBLEM FORMULATION

We offer a formal treatment of the proposed methods here. For our current efforts, we assume two distinct kinds of generative models P: cognitive models P_c and LMs P_{LM}. For a given cognitive model $P_c(y \mid x, \theta)$ with some input x and latent parameters θ, and a language model $P_{\text{LM}}(y \mid x)$, our goal is to find an intervention v from a set of possible interventions $\mathcal{V}$ that minimizes the behavioral divergence between P_c and P_{LM}:

$$v^*(x, \theta) = \arg\min_{v \in \mathcal{V}} D\Big(P_c(\cdot \mid x, \theta) \,\Big\|\, P_{\text{LM}}^{(v)}(\cdot \mid x)\Big)$$

For a chosen P_c, we select a grid of target values $0, ..., k$ the latent θ can take on. For each pair of inputs and cognitive variables (x, θ_k), we sample behavioral outputs $y \sim P_c(\cdot \mid x, \theta_k)$. We then use these samples to form a dataset $\mathcal{D}_{\theta_k} = \{(x, y) \mid y \sim P_c(\cdot \mid x, \theta_k)\}$, which approximates the cognitive model's conditional behavioral distribution at θ_k.

3.2 ACTIVATION AND WEIGHT SPACE INTERVENTIONS

For our current methods, we consider seven interventions in $\mathcal{V}$ across prompts, activations, and weights. Two are prompting baselines: a hard (system) prompt that describes the two extremes of each parameter in prose, and a soft prompt of $k = 8$ trained embedding vectors prepended to

the input (Lester et al., 2021; Li & Liang, 2021). Four intervene on activations, adding a steering vector v to the residual stream at layer ℓ with strength α, $h_\ell \leftarrow h_\ell + \alpha v$. Two of these are standard contrastive baselines that derive v from the options the cognitive model favors at the two ends of a parameter's range, either as their mean activation difference (Rimsky et al., 2024) or by preference optimization (Cao et al., 2024). The other two fit v to the cognitive model's full choice distribution at a target θ: an optimized vector, trained separately for each target, and an affine map, a single learned map from any parameter value to a vector, so new targets need no retraining (Singh et al., 2024). Finally, distillation intervenes on the weights: we fine-tune a LoRA adapter (Hu et al., 2022) on cognitive model rollouts, with a KL penalty toward the base model, and set the target at inference by prepending the parameter values to the prompt (Gupta et al., 2026, In Press). All trained methods minimize the same objective, the KL divergence from the cognitive model's choice distribution to the LM's; Section C gives each method's formulation and training details.

3.3 Evaluation

We evaluate our seven interventions $\mathcal{V}$ on two 8B parameter, instruction-finetuned LMs: `Llama-3.1-8B-Instruct` and `Qwen3-8B` (thinking mode disabled). Our inputs x are sampled from different domains, and our outputs y are distributions over choices within those domains. To compare behavior between P_{LM} and P_c, we compare the log-probabilities over domain choices from each LM Q_{LM} with the probabilities given over those same choices given by $P(y \mid x, \theta)$. Each input x is a scenario given to the LM as a prompt: the scenario text and the K possible choices. The probabilities over the K options are renormalized to sum to one, forming $Q_{LM}(y \mid x)$.

We fit each method on half of the scenarios, using rollouts at parameter settings spread evenly across each variable's range. On the held-out half, we give each method a target θ and recover $\hat{\theta}$ by fitting the cognitive model to the LM's choice probabilities by maximum likelihood (see Equation (5)).

4 Experimental setting: interpersonal communication

A main contribution of our work is establishing the generalization of cognitive steering beyond the toy domains of the originating cognitive models. We focus on high-stakes social interaction, which requires LMs to understand and respond to fundamental processes of interpersonal communication (Berger & Roloff, 2019) in a human-like way.

4.1 Out-of-distribution evaluation scenarios

We adapt seven scenarios from Liu et al. (2023) designed to capture these processes: `bargaining`, `dating_app`, `marriage_argument`, `plane_crash`, `product_launch`, `sharing_secrets`, and `white_lie` (full prompts and their factorial variants are in Section H.1). Examples include an airline spokesperson addressing the press after a plane crash, a CEO launching a new smartphone, and someone drafting a bio for a dating app. Together they cover persuading, negotiating, self-presentation, self-disclosure, delivering bad news, softening an unwelcome truth, and raising a grievance—behaviors theoretically addressed by the cognitive models in Table 1. For each variant, we sample ten generations (temperature 0.8) from the unsteered LM and ten from every steered condition, i.e., each steering method at each target θ. Vector and affine steering use $\alpha = 0.35$ at layer 14 in Llama and $\alpha = 1$ at layer 10 in Qwen.

4.2 Scoring cognitive variables in free generation

Since the cognitive models we explore do not support inference over natural language, we cannot use the originating model to infer the values of the parameters θ from the free generation responses to our OOD scenarios. We instead employ an LLM-as-judge to score the qualities associated with each cognitive parameter. We instantiate this judge using `claude-sonnet-5` at its default settings with reasoning disabled. The judge is provided with a definition of the target θ, a description of what a higher score on the quality looks like, and extraneous features of a generation that should be ignored during evaluation (see Section G.2 for full prompt). The judge is instructed to select which of two generations scores higher on θ, with a "TIE" option.

Model	Unsteered	Prompt	Distilled	Soft prompt	Vector (optimized)	Affine map	Contrastive (optimized)	Contrastive (mean diff.)
IMPR-03	0.13	0.11	0.022 ±0.0017	0.087 ±0.033	**0.00083** ±0.00023	0.021 ±0.0057	0.012 ±0.0053	0.10
LANG-02	0.32	0.28	0.20 ±0.030	**0.19** ±0.017	0.21 ±0.0070	**0.19** ±0.017	0.33 ±0	0.33
PED-02	0.38	0.38	0.018 ±0.0013	0.37 ±0.0065	0.0093 ±0.0034	**0.0057** ±0.0015	0.38 ±0	0.38
RSA-02	0.16	0.15	**0.0063** ±0.00046	0.077 ±0.023	0.011 ±0.0062	0.030 ±0.025	0.12 ±0.014	0.12
RSA-03	0.13	0.089	0.025 ±0.0035	0.12 ±0.029	**0.0025** ±0.0015	0.056 ±0.071	0.15 ±0.0075	0.14
RSA-07	0.18	0.17	0.066 ±0.0037	0.25 ±0.063	0.063 ±0.017	**0.056** ±0.017	0.30 ±0.0022	0.33
RSA-08	0.18	0.21	0.061 ±0.0051	0.15 ±0.073	**0.028** ±0.0060	0.041 ±0.0032	0.13 ±0.033	0.19
RSA-09	0.25	0.25	0.016 ±0.0014	**0.00034** ±0.00024	0.0092 ±0.0039	0.0038 ±0.0038	0.050 ±0	0.041
RSA-11	0.27	0.056	0.094 ±0	**0.027** ±0.012	0.068 ±0.013	0.039 ±0.030	0.053 ±0	0.056
Mean	0.22	0.19	0.057**	0.14	0.045**	0.049**	0.17	0.19
p vs. unsteered	–	0.56	0.0022	0.23	0.0015	0.0025	0.56	0.56

Table 2: **In-distribution parameter recovery in Llama-3.1-8B-Instruct.** Each cell is the MSE between target and recovered parameters θ; the lowest error in each row is in bold. Trained methods show the mean $\pm$ standard deviation over three training seeds. Methods that aren't trained have no seed error. The bottom rows give each method's mean over the nine cognitive models and a paired t-test of its per-model MSE against the unsteered LM ($df = 8$, Holm-corrected over seven comparisons; $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$).

We compare each steered generation with unsteered generations of the same scenario variant, and judge every pair twice, once in each presentation order. We also compare generations steered to different targets of the same parameter, and pairs of two unsteered generations. We fit all three kinds of comparison jointly with a Bradley–Terry–Davidson model (**?**), separately for each cognitive model, parameter, and steering method, with the unsteered LM fixed at 0. This gives a strength for each target θ: the log odds that a response steered toward θ is judged higher than an unsteered response. The slope β of these strengths against the target summarizes how well a method controls the parameter. Unsteered pairs differ by zero by construction, so their fitted strength measures the judge's own noise. Intervals come from resampling scenario variants.

5 Results

5.1 In-distribution steering

We first ask whether each method can steer the LM within the native domains of the cognitive models. For each target setting θ, we recover $\hat{\theta}$ from the LM's response distributions on held-out scenarios (Equation (5)) and measure the mean-squared error (MSE) between the two. We normalize each parameter by its evaluated range before computing error, so that errors are comparable across parameters. For trained methods, we average errors within each training seed and report the mean and standard deviation across three seeds. We compare methods with paired t-tests over the nine cognitive models and Holm-correct all p-values for multiple comparisons.

Three methods stand out in Table 2: optimized vectors, affine maps, and distillation lower recovery error for every one of the nine cognitive models. The unsteered LM has a mean error of 0.22 across models; optimized vectors reduce it to 0.045, affine maps to 0.049, and distillation to 0.057. Each of these reductions is significant (vector: $t(8) = -6.40$, $p = .0015$; affine map: $t(8) = -5.62$, $p = .0025$; distillation: $t(8) = -5.90$, $p = .0022$). The gains are largest on PED-02, where all three methods reduce error from 0.38 to below 0.02, and on IMPR-03, where the optimized vector reduces it from 0.13 to 0.00083.

Among these three methods, however, we find no clear winner. Which one has the lowest error varies from model to model, and no pair differs significantly in mean error (distillation vs. vector: $t(8) = 2.52$, $p = .11$; distillation vs. affine map: $t(8) = 0.96$, $p = .73$; vector vs. affine map: $t(8) = -0.46$, $p = .73$). The remaining methods do not significantly improve on the unsteered LM (prompt: $t(8) = -1.42$, $p = .56$; soft prompt: $t(8) = -2.23$, $p = .23$; contrastive (optimized): $t(8) = -1.44$, $p = .56$; contrastive (mean diff.): $t(8) = -0.89$, $p = .56$). Going forward, we report the OOD generalization performance of the steering methods that performed best in-distribution: distillation, affine map, and vector (optimized).

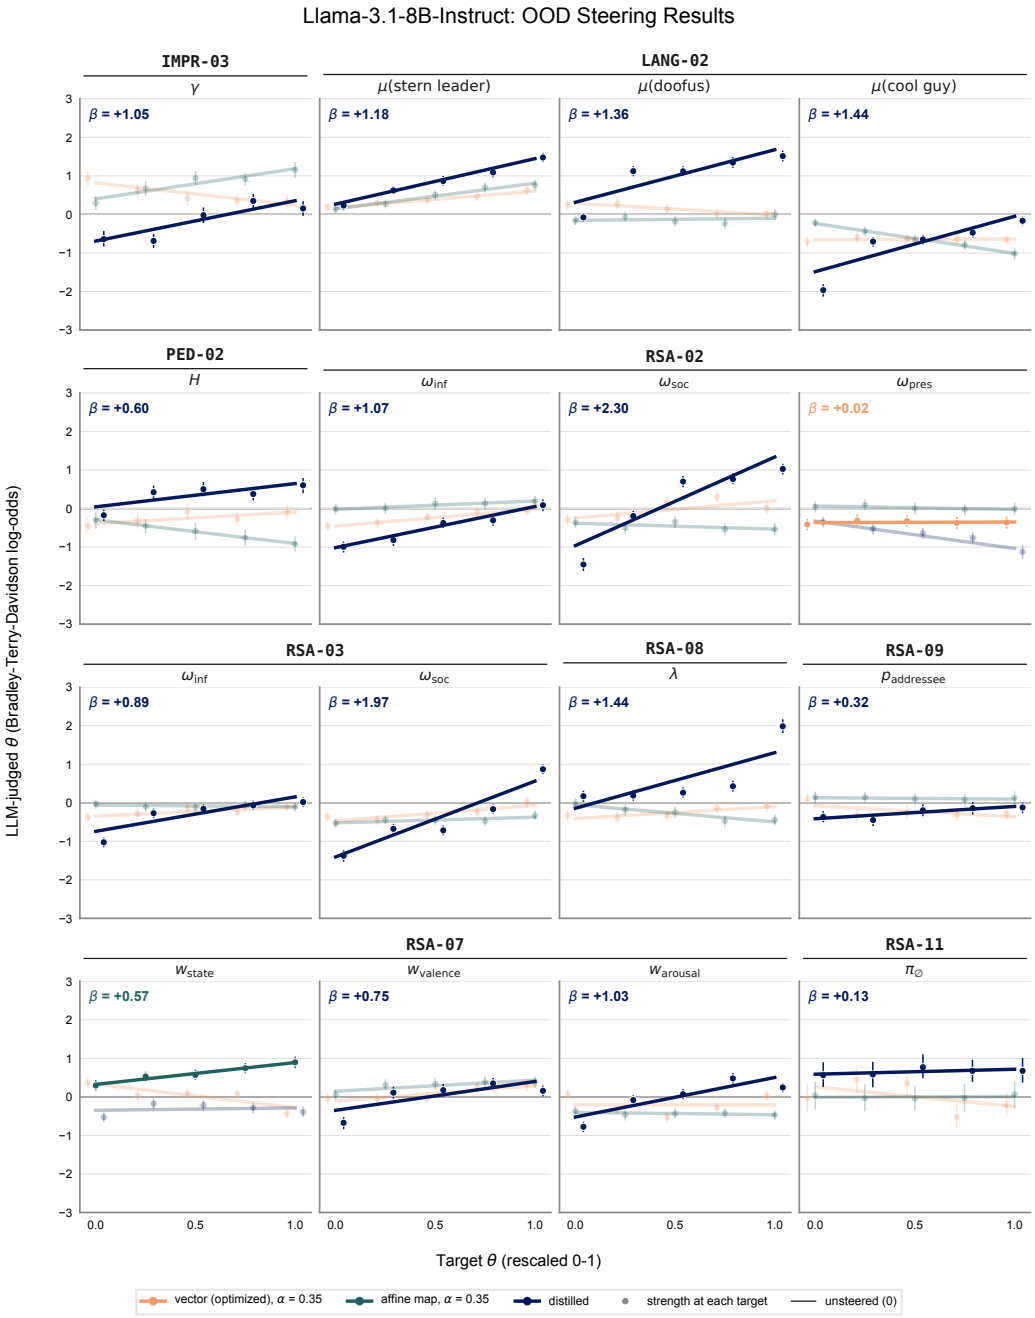

Figure 2: **Out-of-distribution steering results for Llama-3.1-8B-Instruct.** One panel per steered parameter, grouped by cognitive model. The x-axis is the target value of θ, rescaled to 0–1; the y-axis is the LLM-judged θ, in Bradley-Terry-Davidson log-odds, from a joint fit over all pairwise comparisons (steered vs. unsteered, steered vs. steered and unsteered vs. unsteered) with the unsteered model fixed at 0 (judge prompt in Section G.2). Points are the fitted strength at each target with 95% intervals; lines are each method's fitted dose-response slope. Methods: vector (optimized) and affine map steering at $\alpha = 0.35$, and the distilled model. The method with the steepest slope is drawn solid and its slope β printed in the panel; the others are faded.

5.2 Out-of-distribution steering

We next ask whether steering carries over to free generation in the interpersonal scenarios of Section 4. Figure 2 plots, for each parameter, the judged strength of its quality against the target value for Llama-3.1-8B-Instruct (Qwen3-8B results are in Appendix Figure 5); the slope β of this relationship is positive when responses steered toward higher θ are judged to express more of the quality (Section 4.2). Distillation gives the steepest slope for most parameters (14 of 16 in Llama, 11 of 16 in Qwen), and we report the steepest method's slope below.

The clearest control in both LLMs is over the social goals of polite speech. The weight on making the listener feel good, ω_{soc} (Table 1), has the largest slopes of any parameter, both in RSA-02 (Yoon et al., 2020) (Llama: $\beta = 2.30$; Qwen: $\beta = 2.37$) and in RSA-03 (Carcassi & Franke, 2023) (Llama: $\beta = 1.97$; Qwen: $\beta = 3.47$). Both LLMs also respond strongly to RSA-08's trade-off between being relevant and being truthful, λ (Sumers et al., 2024) (Llama: $\beta = 1.44$; Qwen: $\beta = 1.52$), to the value of being seen as friendly in LANG-02 (Burnett, 2019), whether as a cool guy (Llama: $\beta = 1.44$; Qwen: $\beta = 1.58$) or a doofus (Llama: $\beta = 1.36$; Qwen: $\beta = 2.25$), and to the weight on being informative, ω_{inf}, in RSA-02 (Llama: $\beta = 1.07$; Qwen: $\beta = 1.05$) and RSA-03 (Llama: $\beta = 0.89$; Qwen: $\beta = 1.30$). Appendix Table 3 shows Llama generations at the lowest and highest target of several of these parameters.

Other parameters are harder to steer. The prior inclination to say nothing, $\pi_{\varnothing}$ (RSA-11), has no slope distinguishable from zero in either LLM (Llama: $\beta = 0.13$, 95% CI $[-0.26, 0.50]$; Qwen: $\beta = 0.40$, 95% CI $[-0.26, 1.09]$). RSA-02's ω_{pres}, being seen as both honest and kind, is weak (Llama: $\beta = 0.02$, 95% CI $[-0.13, 0.18]$; Qwen: $\beta = 0.33$), and distillation reverses it in both LLMs (Llama: $\beta = -0.71$; Qwen: $\beta = -0.85$). The stern-leader persona is where the two LLMs disagree most: distillation steers it in Llama ($\beta = 1.18$) but reverses it in Qwen ($\beta = -0.52$).

These differences suggest an interpretable basis for which cognitive variables an LLM can take on through steering. The trade-off between informing the listener and making them feel good is fundamental to how people speak politely, and both LLMs track it closely. A prior inclination to stay silent, by contrast, has no natural counterpart in an LLM, which responds to every prompt, and neither LLM can be steered along it. A flat slope could also reflect our scenarios or the judge, but the contrast points to a use for cognitive steering beyond control: as a way to diagnose which of the variables that organize people's behavior are shared with an LLM's representation space.

5.3 Case study: disclosure of negative facts

The judge tells us whether a response reads as expressing a quality, but not what the speaker chooses to say. Two scenarios let us measure a concrete choice instead. In `dating_app`, the speaker writes a bio knowing three true facts about themselves, one of them negative; in `sharing_secrets`, they could tell a new coworker one of three secrets, one of them personal. A separate judge tags which facts each response discloses (Section G.2.1), and Figure 3 shows how often the negative fact is disclosed as each parameter is steered (Appendix Figure 6 shows the same for the personal secret). We report Δ, the change in disclosure from $\theta = 0$ to 1 in percentage points, for the method whose disclosure changes most.

In the dating app, the two informativeness weights behave alike: raising ω_{inf} in either RSA-02 or RSA-03 makes both LLMs more likely to disclose the negative fact (RSA-02: Llama $\Delta = +27$, Qwen $\Delta = +13$; RSA-03: Llama $\Delta = +34$, Qwen $\Delta = +32$), as we would expect of a speaker who wants to be accurate and informative. Raising ω_{soc} has the opposite effect (RSA-02: Llama $\Delta = -26$, Qwen $\Delta = -23$; RSA-03: Llama $\Delta = -24$, Qwen $\Delta = -38$), as does favoring relevance over truthfulness in RSA-08 (Llama $\Delta = -43$; Qwen $\Delta = -20$). Some of the largest drops come from the LANG-02 personas. Valuing being seen as a doofus lowers disclosure in both LLMs (Llama $\Delta = -44$; Qwen $\Delta = -39$), as does the stern leader (Llama $\Delta = -32$; Qwen $\Delta = -14$), and the cool guy lowers it in Llama ($\Delta = -49$) but not Qwen ($\Delta = +6$). This fits the intuition that a speaker managing how they come across has reason to leave an unflattering fact out of their bio. Personal secrets show a less consistent picture (Appendix Figure 6): most parameters change Llama's disclosure by 13 points or less, and in Qwen the doofus and stern-leader personas instead make the personal secret more likely to be shared (both $\Delta = +32$).

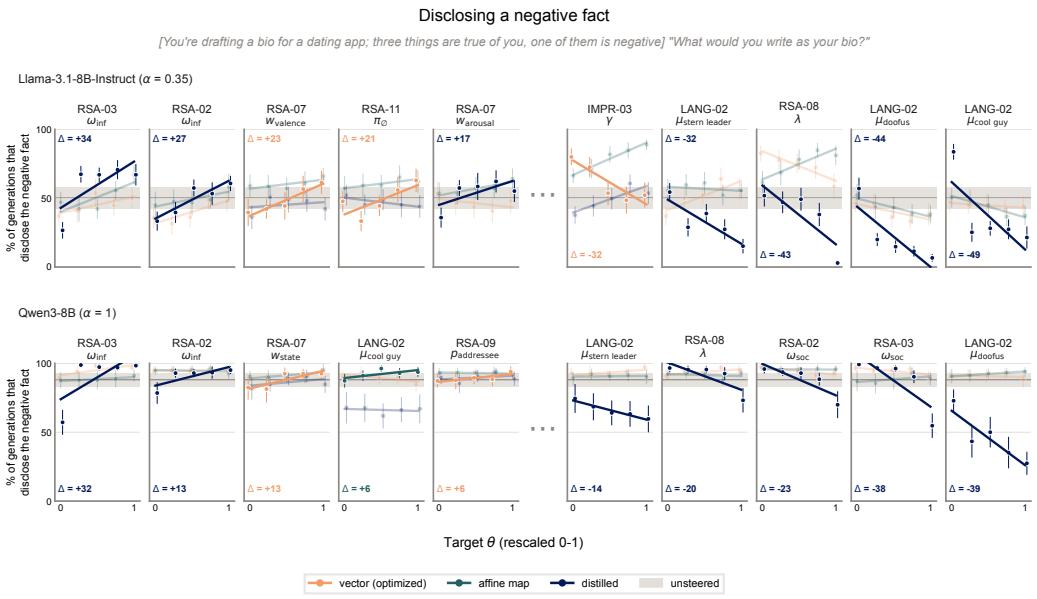

Figure 3: **Out-of-distribution case study: disclosing a negative fact under steering.** Percentage of `dating_app` replies that disclose the speaker's one negative fact (tagged by the disclosure judge, Section G.2.1) against the target θ, for Llama-3.1-8B-Instruct (top) and Qwen3-8B (bottom). Panels show the five parameters with the largest increase in disclosure and the five with the largest decrease. Points are means over scenario variants with 95% intervals; the grey band is the unsteered rate. The method whose rate changes most is drawn solid and labelled with Δ, its change from $\theta = 0$ to 1 in percentage points.

6 Conclusion

We introduced a family of intervention techniques using cognitive models as structured sources of guidance for steering language models via graded dials of behavior that have been tested and verified against real human behavior. We showed that in the native distributions of this family of cognitive models, we are able to reliably increase LM alignment via fine-tuned steering and distillation interventions, and that there are promising signs of out of distribution transfer to naturalistic free-generation scenarios that involve delicate interpersonal dynamics.

Several limitations constrain these conclusions. The model selection process includes LM-based search and an under-powered human validation stage, so the resulting set is by no means exhaustive: further work can easily include a wider set of cognitive models using our methods. Our out-of-distribution evaluation relies on an LM judge because the cognitive models cannot score natural-language generations directly. Although paired presentation orders, ties, and unsteered controls help characterize judge effects, they do not replace human evaluation. The cognitive models also inherit the assumptions and limitations of the human data and experimental tasks from which they were developed. While we test our steering interventions across more than one model family (Llama and Qwen) our results are limited to a single size within this, and should be extended to more model types and sizes in the future. Finally, we want to clarify that successful cognitive steering does not necessarily mean that the LM uses the same cognitive process as a person, only that the intervention produces behavior organized by the same formal variables, establishing an interpretable map from human to LLM behavior. Future work should evaluate these interventions with scaled human judgments, additional model families and domains, and mechanistic tests that tease apart behavioral correlations and shared circuitry.

AI USE STATEMENT

In this work, we used generative AI tools to generate synthetic datasets, implement methods, clean and reformat dataset, support qualitative and thematic data analysis, and interpret results. We have not used generative AI tools for helping develop theoretical models or conceptual frameworks, assisting with translation, or designing or providing feedback on research methodology or experiments, and formulating mathematical claims, providing critical ingredients for proving mathematical claims, assisting in the writing of proofs are not applicable to this work. Additionally, we used generative AI tools for formulating questions for surveys, creating scientific figures or images, transcribing recordings of research material, summarizing existing literature, identifying relevant literature. We have reviewed all AI-assisted work. LLM-generated code was verified and tested for correctness by 1 authors, all LLM-guided literature searches and reviews were monitored, guided, and verified by 2 authors for source legitimacy, all LLM-generated plots and experimental code were audited for correctness. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Asya Achimova, Gregory Scontras, and Martin V. Butz. The alignment model of indirect communication. *PLOS ONE*, 2025. doi: 10.1371/journal.pone.0323839.

Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9798331314385.

Charles R. Berger and Michael E. Roloff. Interpersonal communication. In Don W. Stacks, Michael B. Salwen, and Kristen C. Eichhorn (eds.), *An Integrated Approach to Communication Theory and Research*, chapter 24, pp. 277–292. Routledge, New York, NY, 3 edition, 2019. ISBN 9781138561472. doi: 10.4324/9780203710753-24. URL https://doi.org/10.4324/9780203710753-24.

Eric Bigelow, Daniel Wurgaft, YingQiao Wang, Noah Goodman, Tomer Ullman, Hidenori Tanaka, and Ekdeep Singh Lubana. Belief dynamics reveal the dual nature of in-context learning and activation steering, 2026. URL https://arxiv.org/abs/2511.00617.

Eric J Bigelow, Ekdeep Singh Lubana, Robert P. Dick, Hidenori Tanaka, and Tomer Ullman. In-context learning dynamics with random binary sequences. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=62K7mALO2q.

Egon Brunswik. *The Conceptual Framework of Psychology*, volume 1 of *International Encyclopedia of Unified Science*. University of Chicago Press, 1952.

Heather Burnett. Signalling games, sociolinguistic variation and the construction of style. *Linguistics and Philosophy*, 42(5):419–450, 2019. doi: 10.1007/s10988-018-9254-y.

Yuanpu Cao, Tianrong Zhang, Bochuan Cao, Ziyi Yin, Lu Lin, Fenglong Ma, and Jinghui Chen. Personalized steering of large language models: Versatile steering vectors through bi-directional preference optimization. In *Advances in Neural Information Processing Systems*, 2024.

Fausto Carcassi and Michael Franke. How to handle the truth: A model of politeness as strategic truth-stretching. In *Proceedings of the 45th Annual Conference of the Cognitive Science Society*, pp. 222–228, 2023.

Runjin Chen, Andy Arditi, Henry Sleight, Owain Evans, and Jack Lindsey. Persona vectors: Monitoring and controlling character traits in language models. *arXiv preprint arXiv:2507.21509*, 2025.

Arthur Conmy, Augustine N. Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. Towards automated circuit discovery for mechanistic interpretability. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. URL https://arxiv.org/abs/2304.14997. NeurIPS 2023 Spotlight.

Kenneth J. W. Craik. *The Nature of Explanation*. Cambridge University Press, 1943. ISBN 9780521094450.

Robyn M. Dawes. The robust beauty of improper linear models in decision making. *American Psychologist*, 34(7):571–582, 1979. doi: 10.1037/0003-066X.34.7.571.

Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah D. Goodman. Finding alignments between interpretable causal variables and distributed neural representations. In *Proceedings of the Third Conference on Causal Learning and Reasoning*, volume 236 of *Proceedings of Machine Learning Research*, pp. 421–448. PMLR, 2024. URL https://proceedings.mlr.press/v236/geiger24a.html.

Dedre Gentner and Albert L. Stevens (eds.). *Mental Models*. Psychology Press, 1983.

Gerd Gigerenzer and Daniel G. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669, 1996. doi: 10.1037/0033-295X.103.4.650.

Max Gupta, Changho Shin, Brenden Lake, and Tammy Kwan. Modeling student arithmetic reasoning by distilling cognitive priors into language models, 2026, In Press.

Trygve Haavelmo. The probability approach in econometrics. *Econometrica*, 12(Supplement):iii–vi, 1–115, 1944. doi: 10.2307/1906935.

Edward J. Hu, Yalu Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*, 2022. URL https://openreview.net.

Leonid Hurwicz. On the structural form of interdependent systems. In Ernest Nagel, Patrick Suppes, and Alfred Tarski (eds.), *Logic, Methodology and Philosophy of Science*, pp. 232–239. Stanford University Press, Stanford, CA, 1962.

Philip N. Johnson-Laird. Mental models in cognitive science. *Cognitive Science*, 4(1):71–115, 1980. doi: 10.1207/S15516709COG0401_4.

Justine T. Kao and Noah D. Goodman. Let's talk (ironically) about the weather: Modeling verbal irony. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, pp. 1051–1056, 2015.

Alexander Y. Ku, Declan Campbell, Xuechunzi Bai, Jiayi Geng, Ryan Liu, Raja Marjieh, R. Thomas McCoy, Andrew Nam, Ilia Sucholutsky, Veniamin Veselovsky, Liyi Zhang, Jian-Qiao Zhu, and Thomas L. Griffiths. Levels of analysis for large language models, 2025. URL https://arxiv.org/abs/2503.13401.

Brenden M. Lake, Akshay K. Jagadish, and Guangyuan Jiang. More accurate behavioral predictions with hybrid bayesian-connectionist models. *arXiv preprint arXiv:2608.22154*, 2026. doi: 10.48550/arXiv.2608.22154.

Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021.

Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. *arXiv preprint arXiv:2306.03341*, 2023.

Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, 2021.

Ryan Liu, Howard Yen, Raja Marjieh, Thomas L. Griffiths, and Ranjay Krishna. Improving interpersonal communication by simulating audiences with language models, 2023. URL https://arxiv.org/abs/2311.00687.

Christina Lu, Jack Gallagher, Jonathan Michala, Kyle Fish, and Jack Lindsey. The assistant axis: Situating and stabilizing the default persona of language models. *arXiv preprint arXiv:2601.10387*, 2026.

Robert E. Lucas, Jr. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46, 1976. doi: 10.1016/S0167-2231(76)80003-6.

Jacob Marschak. Economic measurements for policy and prediction. In William C. Hood and Tjalling C. Koopmans (eds.), *Studies in Econometric Method*, number 14 in Cowles Commission Monograph, pp. 1–26. John Wiley & Sons, New York, 1953.

Douglas L. Medin and Marguerite M. Schaffer. Context theory of classification learning. *Psychological Review*, 85(3):207–238, 1978. doi: 10.1037/0033-295X.85.3.207.

Mindaugas Mozuraitis, Suzanne Stevenson, and Daphna Heller. Modeling reference production as the probabilistic combination of multiple perspectives. *Cognitive Science*, 42(S4):974–1008, 2018. doi: 10.1111/cogs.12582.

Aaron Mueller, Jannik Brinkmann, Millicent Li, Samuel Marks, Koyena Pal, Nikhil Prakash, Can Rager, Aruna Sankaranarayanan, Arnab Sen Sharma, Jiuding Sun, Eric Todd, David Bau, and Yonatan Belinkov. The quest for the right mediator: Surveying mechanistic interpretability through the lens of causal mediation analysis, 2025. URL https://arxiv.org/abs/2408.01416.

Sonia K. Murthy, Rosie Zhao, Jennifer Hu, Sham Kakade, Markus Wulfmeier, Peng Qian, and Tomer Ullman. Using cognitive models to reveal value trade-offs in language models. *arXiv preprint arXiv:2506.20666*, 2025.

Sonia Krishna Murthy, Rosie Zhao, Jennifer Hu, Sham M. Kakade, Markus Wulfmeier, Peng Qian, and Tomer Ullman. Cognitive models can reveal interpretable value trade-offs in language models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=nM2QhvybwI.

Allen Newell and Herbert A. Simon. *Human Problem Solving*. Prentice-Hall, 1972. ISBN 9780134454030.

Robert M. Nosofsky. Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115(1):39–57, 1986. doi: 10.1037/0096-3445.115.1.39.

John W. Payne, James R. Bettman, and Eric J. Johnson. Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(3):534–552, 1988. doi: 10.1037/0278-7393.14.3.534.

Ciyang Qing, Noah D. Goodman, and Daniel Lassiter. A rational speech-act model of projective content. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pp. 1110–1115, 2016.

Nina Rimsky, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Turner. Steering llama 2 via contrastive activation addition. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15504–15522, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.828. URL https://aclanthology.org/2024.acl-long.828/.

Lance J. Rips. *The Psychology of Proof: Deductive Reasoning in Human Thinking*. MIT Press, 1994. ISBN 9780262181532. doi: 10.7551/MITPRESS/5680.001.0001.

Rachel Ryskin, Suzanne Stevenson, and Daphna Heller. Probabilistic weighting of perspectives in dyadic communication. In *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society*, 2020. URL https://escholarship.org/uc/item/35m9s3fp.

Johannes A. Schubert, Akshay K. Jagadish, Marcel Binz, and Eric Schulz. In-context learning agents are asymmetric belief updaters. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Herbert A. Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69 (1):99–118, 1955. doi: 10.2307/1884852.

Shashwat Singh, Shauli Ravfogel, Jonathan Herzig, Roee Aharoni, Ryan Cotterell, and Ponnurangam Kumaraguru. Representation surgery: Theory and practice of affine steering. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Theodore R. Sumers, Mark K. Ho, Robert D. Hawkins, Karthik Narasimhan, and Thomas L. Griffiths. How to talk so AI will learn: Instructions, descriptions, and autonomy. In *Advances in Neural Information Processing Systems*, volume 35, 2022. doi: 10.52202/068431-2519.

Theodore R. Sumers, Mark K. Ho, Thomas L. Griffiths, and Robert D. Hawkins. Reconciling truthfulness and relevance as epistemic and decision-theoretic utility. *Psychological Review*, 131 (1), 2024.

Daniel Chee Hian Tan, David Chanin, Aengus Lynch, Brooks Paige, Dimitrios Kanoulas, Adrià Garriga-Alonso, and Robert Kirk. Analysing the generalisation and reliability of steering vectors. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024.

Amos Tversky. Features of similarity. *Psychological Review*, 84(4):327–352, 1977. doi: 10.1037/0033-295X.84.4.327.

Tomer D Ullman and Joshua B Tenenbaum. Bayesian models of conceptual development: Learning as building models of the world. *Annual Review of Developmental Psychology*, 2(1):533–558, 2020.

Tomer D. Ullman, Elizabeth Spelke, Peter Battaglia, and Joshua B. Tenenbaum. Mind games: Game engines as an architecture for intuitive physics. *Trends in Cognitive Sciences*, 21(9):649–665, 2017. doi: 10.1016/J.TICS.2017.05.012.

Erica J. Yoon, Michael Henry Tessler, Noah D. Goodman, and Michael C. Frank. Polite speech emerges from competing social goals. *Open Mind*, 4:71–87, 2020. doi: 10.1162/opmi_a_00035.

A COGNITIVE MODEL SELECTION

A.1 AUTOMATED SEARCH CRITERIA

Candidate models were assembled by an automated, agent-driven pipeline that identifies published cognitive models from a literature search over relevant venues. We then screened candidates on functional criteria, which follow from how we evaluate steering: an LM generates a single reply to a single scenario, and an LLM judge reads that reply. Three are hard filters:

- *Stateless*: the parameter can be fit from one generated response to one scenario. A preceding message as context is fine; multiple turns, a convention formed across rounds, or a training run is not
- *Recoverable from the scenario*: the parameter is identifiable from facts a scenario can state (what happened, what the other person said, who can see what, what is at stake). Models that need a distribution only an elicitation experiment produces (e.g., a prior over unobserved worlds or normed ratings) are excluded
- *Dispositional parameter*: the parameter is something the speaker values, believes, or is trying to be, rather than a rationality or precision term, a per-word cost, or a semantic threshold, since the judge rates how much of a quantity a single reply shows

The remaining criteria were weighed together:

- *Interpersonal process*: the model is about one of the ten fundamental processes of interpersonal communication (Berger & Roloff, 2019) (e.g., social support, deception, bargaining and negotiation, impression management), judged against the processes rather than our evaluation scenarios, so that the set is not tuned to the out-of-distribution test
- *Generality*: the construct holds outside the original task, and is not defined by a reference game, a picture array, or one specific word; relatedly, it is not tied to one type of scenario
- *Behavioral level*: the model is about what the speaker is doing, not which word comes out (syntax, lexical choice, phonology)
- *Available model code*: the authors' implementation exists, so that there is a reference to port and check our implementation against
- *Marginal coverage*: as a tie-breaker, a candidate that covers a process no other model covers is preferred over one that duplicates a covered process

B EXAMPLE OF COGNITIVE MODEL ROLLOUTS

Below is an example of $\mathcal{D}_{\theta_k}$ generated by a cognitive model of polite speech (Yoon et al., 2020). Given a scenario x, the model weights three competing utilities, being informative, kind, and projecting a particular balance of goals (presentational), with θ denoting the resulting triplet of weights. x describes a speaker who believes a listener's cookies are terrible but is asked to rate them, and the distribution over responses y depends on the speaker's goals θ[3]. Below, we list four speaker profiles by type, parameter setting θ, and distribution over utterances y for the cookie scenario x:

speaker type	$\boldsymbol{\theta}$	**rollout** {response: probability}
informative	$(1, 0, 0)$	{terrible: 0.79, bad: 0.18, amazing: 0.02, ...}
mixed	$(0.5, 0.25, 0.25)$	{terrible: 0.47, bad: 0.27, amazing: 0.16, good: 0.07, ...}
kind, face-saving	$(0, 0.5, 0.5)$	{good: 0.45, amazing: 0.33, bad: 0.10, terrible: 0.07, ...}
face-saving only	$(0, 0, 1)$	{terrible: 0.35, bad: 0.34, good: 0.14, amazing: 0.07, ...}

C STEERING METHODS

C.1 PROMPTING BASELINES

To highlight the effectiveness of our methods for steering and distillation, we first evaluate P_{LM} on a simple set of hard (system) prompting and soft prompting baselines. In a hard prompting setting,

[3]We omit the model's formal definition to focus on what D_{θ_k} contains.

we describe the setting in prose to P_{LM} and for each variable in θ, we define what its two extreme settings mean via an LM-generated description, which we validate manually.

In soft prompting, following Lester et al. (2021); Li & Liang (2021) we train and prepend $k = 8$ embedding vectors to the input prompt describing the task with the weights of P_{LM} remaining frozen. Training parameters and objectives are the same as described in Equation (2); in this way, we can directly compare interventions in the learned parameters, the prompt, and the residual stream.

C.2 Activation space interventions

Let h_ℓ denote the residual-stream activation at layer ℓ. A steering vector $v \in \mathbb{R}^{d_{\text{model}}}$ is added at every token position,

$$h_\ell \;\leftarrow\; h_\ell + \alpha\, v, \tag{1}$$

where $\alpha \in \mathbb{R}$ is the intervention strength; negative α steers away from the intended behavior. We compare four ways of obtaining v. We sweep over α, ℓ, and token positions to induce the desired behavior and to not break the model's coherence by overfitting to the cognitive model's behavior (see Figure 7).

C.2.1 Contrastive and BiPo

As baselines, we also compare two standard ways of extracting a steering vector from contrastive examples, by generating pairs mined from each cognitive model. For each parameter θ_j, we take the two ends of its range and, for every training scenario x, the option on which the cognitive model places most probability mass at each end, formally we compute the contrastive pair: $y^+ = \arg\max_y P(y \mid x, \theta_{hi})$ and $y^- = \arg\max_y P(y \mid x, \theta_{low})$. Contrastive activation steering (CAA) Rimsky et al. (2024) takes the residual activation at layer ℓ on the answer token and creates steering vector v to be the mean activation difference under y^+ and y^-. Bi-directional preference optimization (BiPo) Cao et al. (2024) goes a step further by training v by gradient descent using the difference between y^+ and y^- as the objective function.

C.2.2 Finetuned steering vectors

A finetuned steering vector uses the full target distribution, optimizing v directly against it,

$$v^\star \;=\; \arg\min_v \sum_x \mathrm{KL}\big(P_\theta(\cdot \mid x) \,\big\|\, Q_v(\cdot \mid x)\big)\,, \tag{2}$$

which reduces to cross-entropy under the target. To finetune a single vector per intended behavioral target uses Adam with 150 steps, learning rate 0.02, batch size 8, and Gaussian initialization ($\sigma = 0.05$); sweeps repeat each target over 3–10 random seeds. Because each vector must be individually finetuned per intended behavioral stance, we now turn our attention to a more efficient way of interpolating across a family of trained steering vectors by learning an affine map.

C.2.3 Affine Steering

Here, we introduce a single plane of *affine* vectors that map *any* cognitive parameter value to a vector. Similar techniques define affine steering maps to and from LM activations Singh et al. (2024), while our map is affine between cognitive parameters and LM activations.

For a simplex-valued parameter,

$$V(W) \;=\; \sum_i \big(a_i\, w_i + b_i\big)\, \hat{v}_i, \tag{3}$$

with one unit direction $\hat{v}_i$ and one slope–intercept pair (a_i, b_i) per component. The parameters $\{\hat{v}_i, a_i, b_i\}$ are trained end-to-end under the same objective (Equation (2)), with stances sampled from the interior of the simplex. Picking a new stance involves a simple linear interpolation on the simplex and a cheap matrix-vector product (Equation (3)), rather than a fresh round of optimization. We train each map with 1,000 steps of AdamW optimization with an identical learning rate and batch size to the single vector fits.

C.3 Weight space interventions: Cognitive Model Distillation

Following prior work on distillation Lake et al. (2026); Gupta et al. (2026, In Press), we introduce a train-time method for updating LM weights using cognitive model rollouts as fine-tuning data. We train a LoRA adapter Hu et al. (2022) by minimizing the cross-entropy between the teacher's (cognitive model's) choice distribution and the student's (LM's) choices. The objective is equivalent to Equation (2), with a KL-anchor to the base model:

$$\mathcal{L}(\phi) \;=\; \underbrace{\mathrm{KL}\big(P_\theta(\cdot \mid x) \,\big\|\, Q_\phi(\cdot \mid x)\big)}_{\text{adapt to cognitive model target}} \;+\; \beta \underbrace{\mathrm{KL}\big(Q_\phi(\cdot \mid x) \,\big\|\, Q_0(\cdot \mid x)\big)}_{\text{regularization}} \tag{4}$$

To recover the desired cognitive model parameterization at inference time, we use cognitive parameter prompting (CPP) Gupta et al. (2026, In Press), where the targeted parameters and their desired values are prepended to a user's prompt. For example, after distilling the polite speech model of Yoon et al. (2020), we can prompt for highly social, but minimally informative and presentational behavior by attaching the prefix `[informational=0.20; social=0.70; presentational=0.10]` to any prompt.

C.4 Recovering $\hat{\theta}$

For evaluation, we split the scenarios into train and test. On the training half, we generate rollouts at parameter θ settings spread evenly across each variable's range and fit every method in Section 3 via the loss objectives between P_c and Q_{LM}. On the evaluation half, we give each method a target setting θ and read off $Q_{LM}(y \mid x)$ for every scenario $x \in \mathcal{X}$. We then fit the cognitive model to the LM's response probabilities via maximum likelihood estimation to obtain $\hat{\theta}$

$$\hat{\theta} \;=\; \arg\max_{\theta} \sum_{x \in \mathcal{X}} \sum_{y=1}^{K} Q(y \mid x)\, \log P(y \mid x, \theta). \tag{5}$$

D In-distribution steering results

D.1 Localizing cognitive variable representations

We use Distributed Alignment Search (DAS) (Geiger et al., 2024) to test whether cognitive steering makes cognitive variables localizable in LM activation space. Here, "setting" an LM to θ means steering it toward that cognitive parameter value using one of our learned interventions: an activation-space intervention $v(\theta)$ or, for a distilled model, cognitive parameter prompting. DAS compares two runs of the same frozen LM. In the source run, the LM is steered toward θ_s; in the base run, it is steered toward a different setting θ_b. At a given layer, DAS replaces the component of the base activation in a learned subspace with the corresponding component from the source activation, while retaining the rest of the base activation.

If the patched base run preserves its context x_b but behaves as though it had been steered toward θ_s, the learned subspace provides evidence for a causal alignment with θ. We therefore compare the patched LM's behavior with the cognitive-model target $P_c(y \mid x_b, \theta_s)$. DAS learns the n-dimensional orthonormal subspace within the $\mathbb{R}^d$ activation space that best reproduces this target behavior.

We train DAS using source runs steered toward the extrema of each θ and evaluate patched base runs on held-out contexts x_b. We measure *variable recovery* as the slope of the recovered $\hat{\theta}$ against the source target θ_s, where 1.0 indicates full recovery and 0 indicates none. Recovery from the unpatched source run provides the reporting ceiling.

As shown in Figure 4, cognitive variables are not reliably recoverable from the unsteered model but become recoverable in middle to late layers after cognitive steering. Distillation with CPP produces stronger and more persistent recovery than affine and finetuned steering vectors. Thus, cognitive steering may create, strengthen, or reorganize representations of the cognitive variables rather than merely reveal representations already present in the base model.

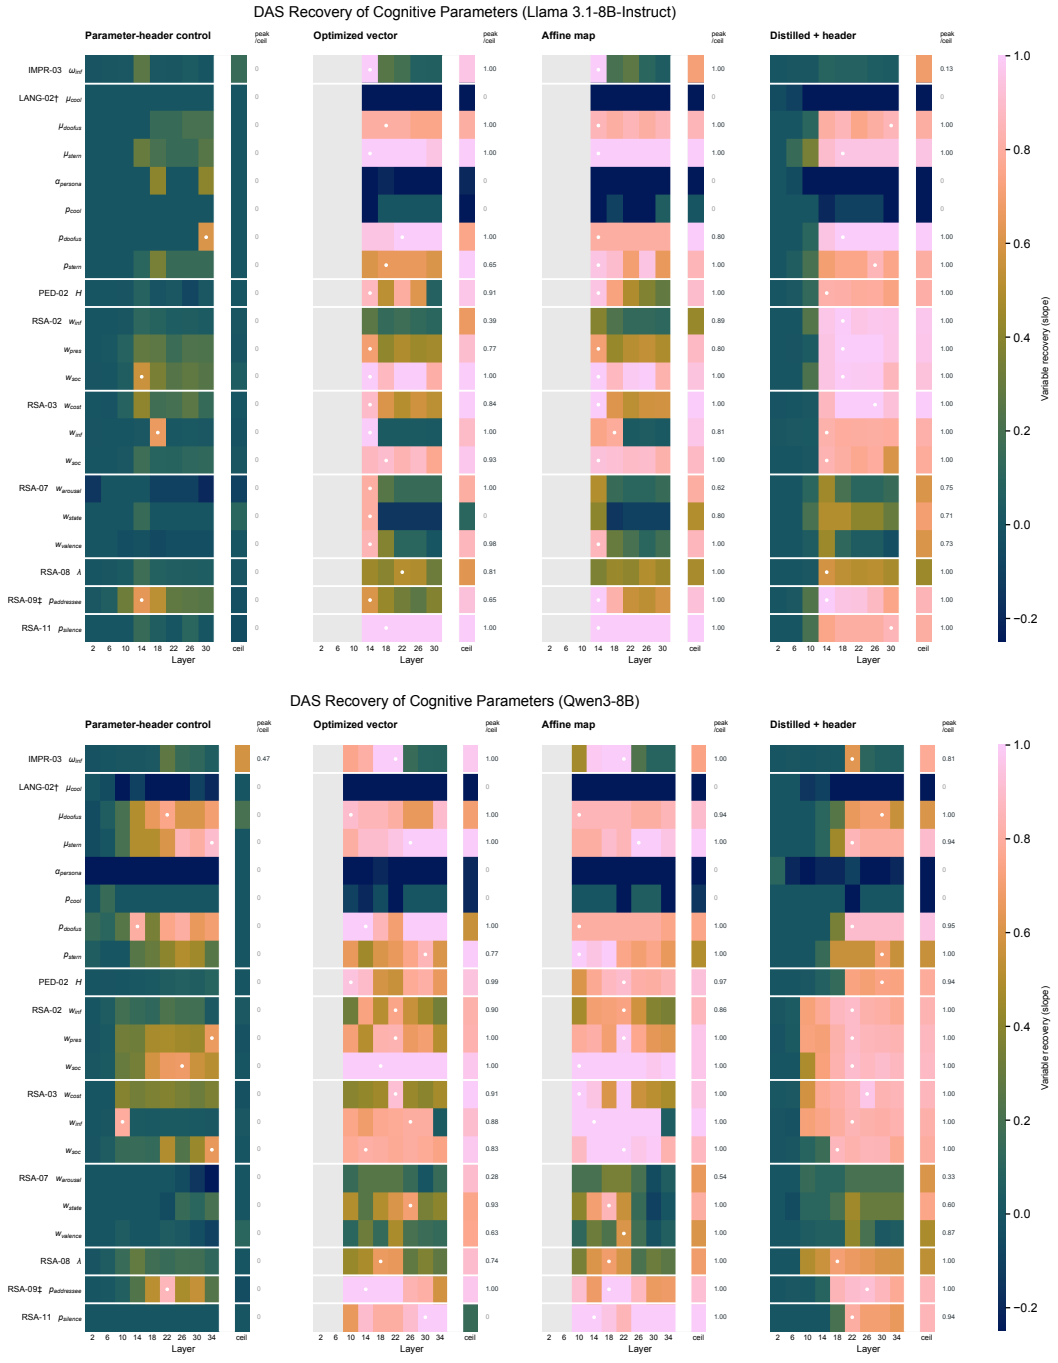

Figure 4: **Recovery of cognitive parameters with distributed alignment search (DAS) in Llama 3.1-8B-Instruct** Rows show cognitive parameters and columns show layers, grouped by steering method. Colour measures how strongly recovered parameters track their targets: 1 indicates unit sensitivity and 0 indicates none. The "ceil" column gives recovery without patching; white dots mark the strongest layer when recovery exceeds 0.5. Grey cells were not evaluated. Results use one training seed and one DAS initialization

E OUT-OF-DISTRIBUTION STEERING RESULTS

Table 3 shows example Llama generations at the lowest and highest target θ. Figure 5 shows the judged strength of each steered parameter against its target for Qwen3-8B, the counterpart to Figure 2 in the main text (Section 4.2). Figure 6 shows disclosure of the personal secret in `sharing_secrets`, the counterpart to Figure 3 in the main text.

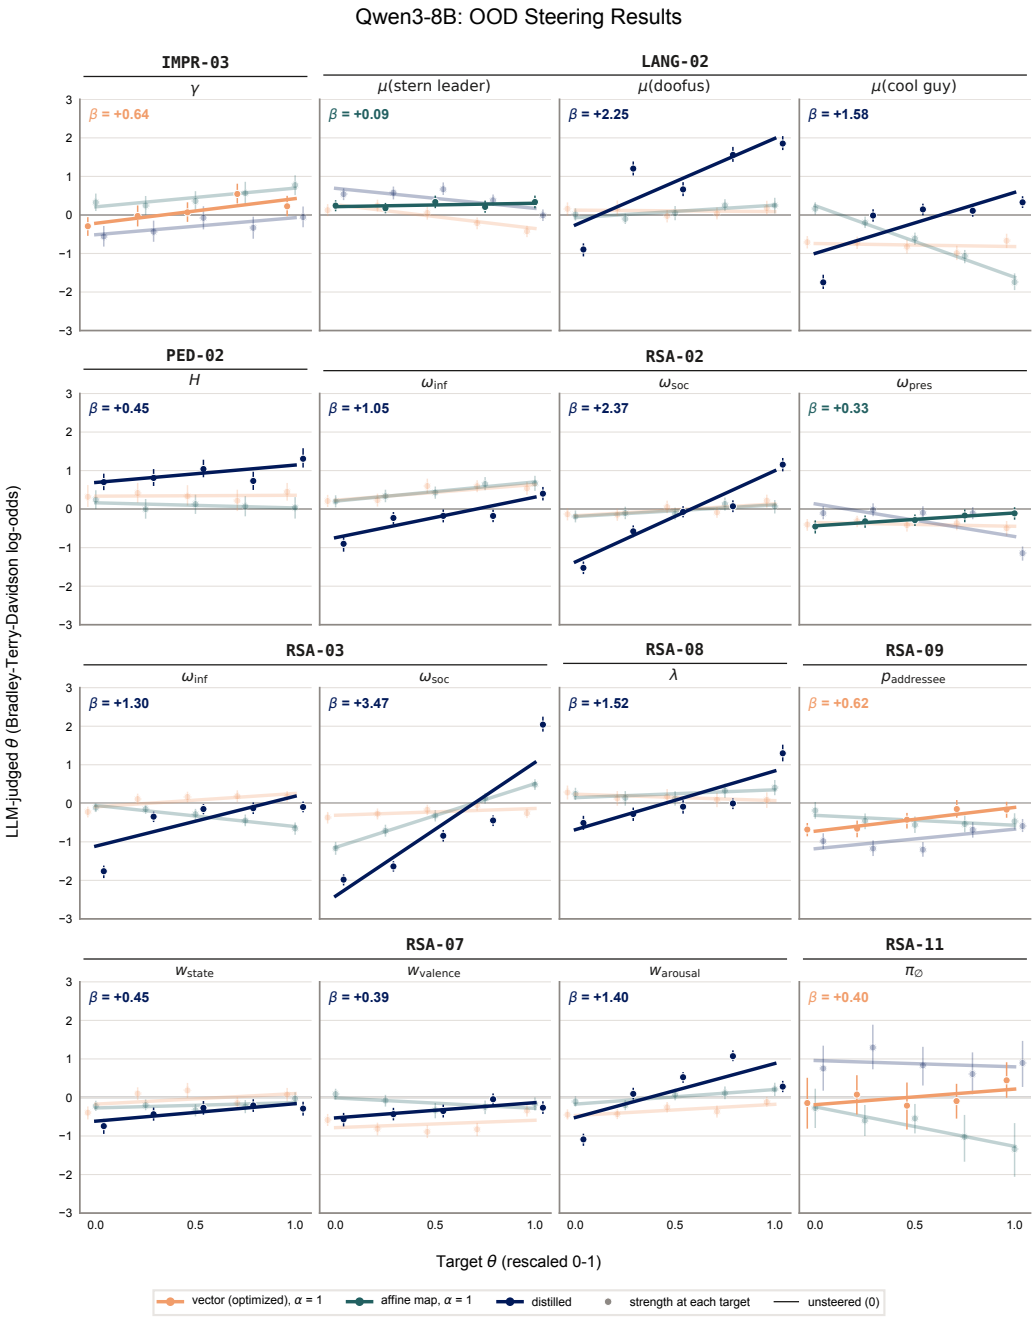

Figure 5: **Out-of-distribution steering results for Qwen3-8B.** Same as Figure 2, but for Qwen3-8B with vector (optimized) and affine map steering at $\alpha = 1$.

Table 3: **Example generations under steering at low and high target θ.** Llama-3.1-8B-Instruct generations at the lowest and highest target θ for five of the most steerable parameters OOD (slope β; Figure 2), all steered by distillation. Each parameter is shown in the scenario where its Bradley–Terry strength, the log odds that a steered generation is judged to express more of the parameter's quality than an unsteered one (Section 4.2), differs most between the two targets. In each pair, the judge rated the highest-target generation higher in both presentation orders.

Parameter	Lowest target	Highest target
RSA-02 ω_{soc} $\beta = 2.30$ `sharing_secrets`	$\omega_{\text{soc}} = 0$*:* I've been meaning to talk to you about something, but I'm not sure why I'm sharing it now. I guess it just feels like we're getting to know each other pretty well. To be honest, I've been feeling pretty anxious about my job lately. I'm not sure if it's just because it's new or what, but I've been feeling a little overwhelmed by it all. It's not something I usually talk about, so don't worry, I'm fine, I just needed to get that off my chest.	$\omega_{\text{soc}} = 1$*:* "Hey, I really like getting to know you, and I feel like we're really clicking. I wanted to share something a bit more personal with you. To be honest, I've been really nervous about starting this new job and I'm still getting used to everything. I'm trying to figure things out and get more confident, but sometimes it feels like I'm taking a big leap of faith. I guess it's nice to have someone to talk to about it, and I value your opinion. What about you? How's it going so far?"
RSA-03 ω_{soc} $\beta = 1.97$ `bargaining`	$\omega_{\text{soc}} = 0$*:* I'd like to ask, is the $340 price set by the artist who made this rug, or is that your price for it?	$\omega_{\text{soc}} = 1$*:* I love the colors and design of this rug! The quality is amazing. I was wondering if you could give me any information about the person who made it? Sometimes I like to support the artisans directly, but I also appreciate the work you do in showcasing their talents here at the market.
RSA-08 λ $\beta = 1.44$ `dating_app`	$\lambda = 0$*:* "Hi, I'm a business major at the university, trying to figure out my career path after graduation. I'm a bit of an outdoorsy person, having recently biked solo across my home state last summer. I'm also a bit of a realist when it comes to enjoying a good drink with friends - I've learned from my mistakes. I'm looking for someone who values honesty and can appreciate a laid-back, adventurous spirit. If you think we could hit it off, let's grab a drink and see where things go."	$\lambda = 1$*:* "Ambitious business major turned world traveler, always on the move and ready for the next adventure. You know those guys who are always up for a spontaneous hike or a night out? That's me. Let's grab a drink and see where the night takes us."
LANG-02 $\mu_{\text{cool guy}}$ $\beta = 1.44$ `marriage_argument`	$\mu_{\text{cool guy}} = 0$*:* I'm really upset that you didn't come home for our anniversary. I know I brought it up last minute and you probably forgot, but that's even more frustrating to me because you're the one who's always reminding me of important dates. It feels like your priorities were elsewhere and that's hurtful.	$\mu_{\text{cool guy}} = 4$*:* "Honey, I wanted to talk to you about last night. I know we've talked about our anniversary before and how much it means to me, but I was really disappointed that you didn't make an effort to come home early and spend some time with me. I understand that you wanted to spend time with your friends, and that's fine, but for me, our anniversary is a special time for us to reconnect and celebrate our love. Can we talk about how we can make this work for both of us next year?"
RSA-02 ω_{inf} $\beta = 1.07$ `dating_app`	$\omega_{\text{inf}} = 0$*:* "Hey there, I'm [name], a business major by day and a musician by night. When I'm not juggling classes or gigs, you can find me trying out new restaurants or exploring the city. I'm a hopeless romantic looking for someone who can appreciate a good melody and a great sense of humor. I'm a bit quirky, a bit goofy, and always up for an adventure. If you can quote a good song, make a decent joke, or just appreciate a good laugh, let's grab coffee and see where things go."	$\omega_{\text{inf}} = 1$*:* "I'm a business major trying to get back on track academically after a rough semester. I'm not perfect, but I'm working on it. When I'm not studying, you can find me shredding on guitar with my band. We're starting to get some paid gigs around town, so that's a fun new adventure. I'm looking for someone who can appreciate a good jam session and isn't too picky about a guy who's still figuring things out. If you're down for a low-key coffee date or concert night, let's connect."



Figure 6: **Disclosing a personal secret under steering.** Same as Figure 3, but for the `sharing_secrets` scenario, in which the speaker could share one of three secrets with a new coworker, one of them personal; the percentage of replies that disclose the personal secret.

F Generation quality under steering

F.1 Knowledge degradation

Steering more strongly (with higher magnitudes α) is often shown to degrade text coherence (by introducing repetition, looping or just random token generation) and general knowledge. In the coherence plots below, we highlight coherence remains stable at a given fixed steering magnitude. Here we show how we arrived at that magnitude by sweeping across strengths and layers, and observing how in-distribution steering MSE (Figure 7, left) and GSM8K and MMLU accuracy (Figure 7, middle and right, right axis) degrade on average across all 9 cognitive models. We arrive at $\alpha = 0.35$ as an optimal tradeoff between low MSE, high benchmark accuracy, and high textual coherence so that steered outputs remain factually correct, coherent, and indicative of the cognitive model's intended behavior.

F.2 Textual coherence

Figures 8 and 9 show the share of coherent, degraded, and incoherent replies under each steering method for Llama and Qwen.

G LLM-as-Judge Prompts

We use `claude-sonnet-5` as a judge at several steps of the OOD evaluation pipeline. We start with a basic filter for degraded and incoherent outputs: the judge rates each reply as coherent, degraded or incoherent, and incoherent replies are excluded before any comparison is made (Section G.1). All of our steering methods pass this filter with high rates of coherent replies (Figures 8 and 9). Then, we determine the degree to which the OOD generations demonstrate the qualities illustrated by the cognitive variable used as the target for steering, using a pairwise judge (Section G.2). Finally, we implement a scenario-specific judge for analyzing the disclosure rate of responses to particular scenarios whose design allows the responder to choose which of three facts – that vary in severity or personal implication – they disclose about themselves, (Section G.2.1).

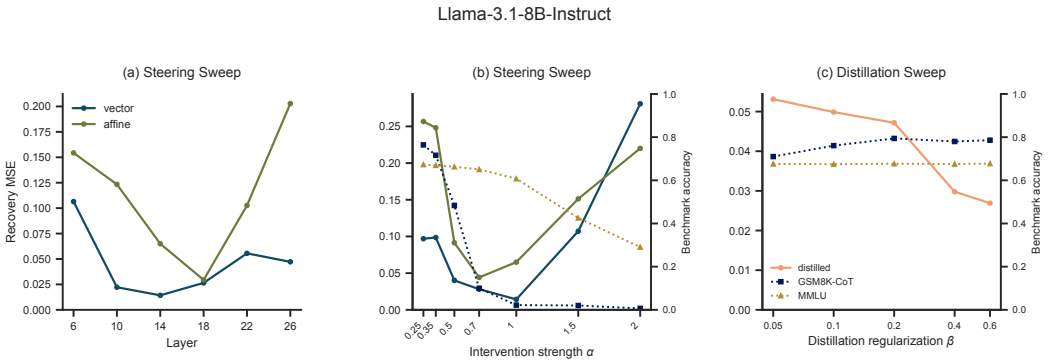

Figure 7: **Steering and distillation sweeps for Llama-3.1-8B-Instruct across nine cognitive models.** (a) Parameter-recovery MSE as a function of intervention layer at steering strength $\alpha = 1$. (b) Recovery MSE as a function of steering strength at layer 14, with GSM8K and MMLU accuracy under optimized-vector steering shown on the right axis. (c) Recovery MSE and benchmark accuracy as a function of the base-model KL regularization coefficient β during distillation. Solid lines show recovery MSE; dotted lines show benchmark accuracy. Curves report medians across cognitive models, after averaging over parameter axes within each model, using one fitted seed. Lower MSE indicates better parameter recovery; higher accuracy indicates better benchmark performance.

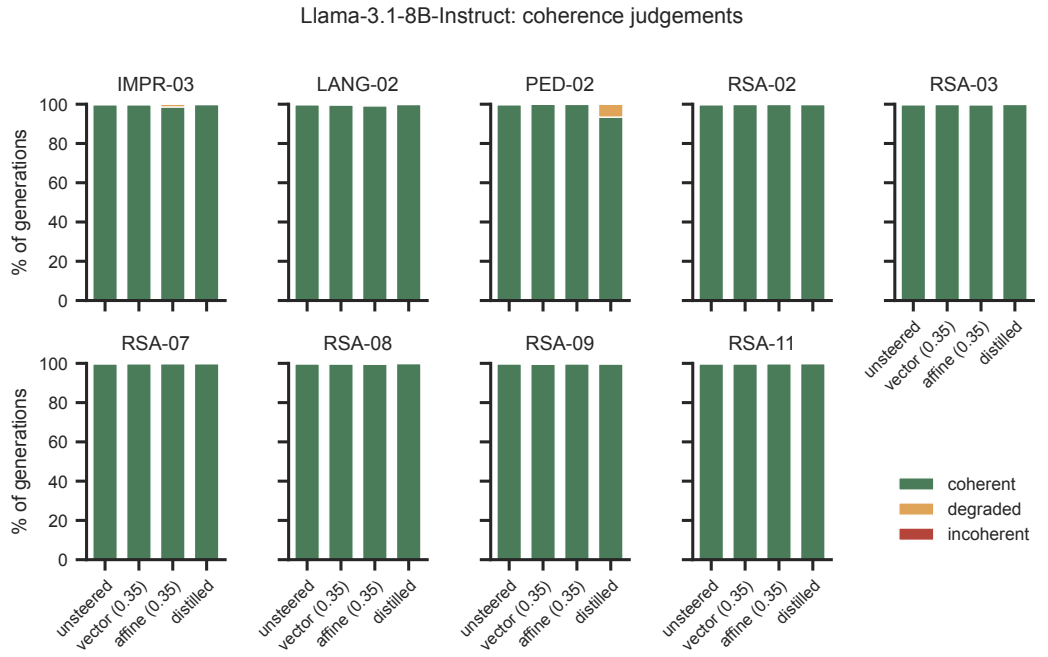

Figure 8: **Coherence of Llama-3.1-8B-Instruct generations.** Share of replies to the out-of-distribution scenarios rated coherent, degraded or incoherent by the coherence judge (prompt in Section G), one panel per cognitive model, for unsteered generation, vector and affine steering (magnitude 0.35, layer 14) and the distilled model.

G.1 Coherence of generations under steering

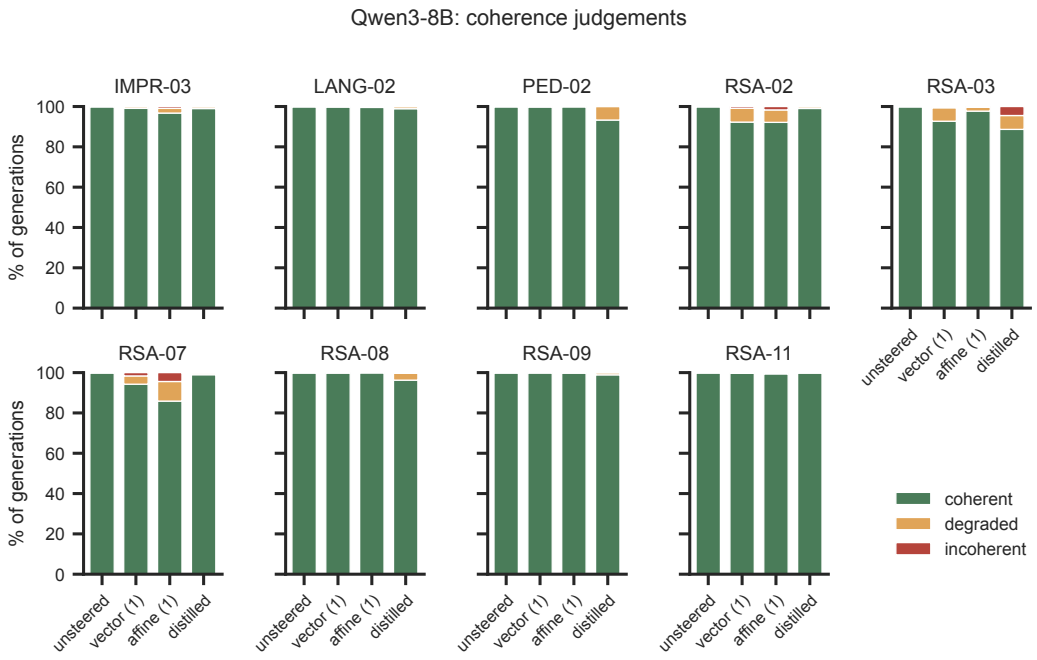

Figure 9: **Coherence of Qwen3-8B generations.** Same as Figure 8, but for Qwen3-8B with vector and affine steering at magnitude 1.0, layer 10.

Coherence judge

```
You are checking whether a piece of text is well-formed language. A language
model was given the scenario and question below and wrote the reply. Some
replies come from a model whose internals were being modified, and those
sometimes break down into repetition, fragments or garbage. Your job is to say
how intact the reply is AS LANGUAGE.

Rate ONLY the form of the reply. Do NOT judge whether it is a good, helpful,
accurate, polite, appropriate or on-topic answer, whether it follows the
question's instructions, or its tone, style or length. A fluent reply that
ignores the question, invents details, contradicts the scenario or says
something strange is still COHERENT. The scenario is shown only so you can
read the reply in context.

COHERENT: reads as fluent text a person could have written. Sentences are
grammatical and connect to each other. Minor slips are fine, and so is a short
reply: one complete sentence is COHERENT.
DEGRADED: still recognisably an attempt at a reply, but visibly broken in
form. For example a phrase or sentence repeated, text that trails off or is
cut off mid-thought, a stray fragment or garbage appended to otherwise fluent
text, or clauses that do not connect.
INCOHERENT: not functioning as language. For example word salad, strings of
symbols, numbers or code-like tokens, a loop that takes over most of the text,
or a bare heading or fragment with no content.

SCENARIO
SCENARIO TEXT

WHAT THE RESPONDER WAS ASKED
QUESTION
```

```
REPLY
REPLY

How intact is the reply as language? Answer with exactly one word: COHERENT,
DEGRADED, or INCOHERENT.
```

G.2 COGNITIVE VARIABLE SCORING

Pairwise judge

```
You are comparing two responses to the same scenario, written by two different
speakers.
Your task is to judge which response scores HIGHER on ONE specific dimension:
[DIMENSION NAME].

A response scores HIGHER on this dimension if it [WHAT SCORES HIGHER].

Judge ONLY this dimension -- not helpfulness, writing quality, length, or
correctness. Ignore [WHAT TO IGNORE].
If one response is even slightly further toward the HIGHER end, pick that one.
Answer TIE only when neither response is further toward the HIGHER end than
the other; differences in wording, detail, or length alone do not make one
response score higher. Answer TIE if the dimension does not apply to this
scenario.

SCENARIO
SCENARIO TEXT

WHAT THE RESPONDER WAS ASKED
QUESTION

RESPONSE A
RESPONSE A

RESPONSE B
RESPONSE B

Which response scores higher on the stated dimension? Answer with exactly one
word: A, B, or TIE.
```

RSA-02, ω_{inf}

```
Dimension name: informational utility (be understood)

Scores higher if it: lets the reader work out the actual situation described
in the scenario, even where that is unwelcome. In the model this is weighed
against ω_soc (pleasing the listener) and ω_pres (projecting a stance)

Ignore: whether the response is kind, well-written, or appropriate
```

RSA-02, ω_{soc}

```
Dimension name: social utility (make the listener feel good)

Scores higher if it: says what will please the person being addressed, whether
or not it is true. In the model this is weighed against ω_inf (conveying the
true state) and ω_pres
```

Ignore: whether the response is accurate, or well-intentioned

RSA-02, ω_{pres}

Dimension name: presentational goal: being seen as someone with a particular balance between being informative and being kind

Scores higher if it: is worded so that the speaker comes across as someone who [BALANCE]. Judge how strongly the reply presents the speaker this way, beyond what its message says about the situation

[BALANCE], by the ϕ the replies were made with (ϕ = 0.5 in the main runs):
ϕ = 0: cares mainly about being informative: about telling the listener how things really are
ϕ = 0.5: cares about being informative and about being kind in roughly equal measure: about telling the listener how things really are while sparing their feelings
ϕ = 0.95: cares mainly about being kind: about sparing the listener's feelings

Ignore: whether that is a good balance to have, whether the reply is accurate, and how well it is written

RSA-03, ω_{soc}

Dimension name: expression weight on social value (choose kind wordings)

Scores higher if it: chooses phrasing whose plain, literal meaning is the most flattering available. In the model this is weighed against ω_{inf} and ω_{cost}

Ignore: whether that flattering reading is accurate

RSA-03, ω_{inf}

Dimension name: expression weight on informativity (be understood)

Scores higher if it: states its position precisely, so a reader would land on exactly the state of things the speaker meant, with little room for misreading. In the model this is weighed against ω_{soc} and ω_{cost}

Ignore: whether the position taken is TRUE -- a confident, specific falsehood scores high here and a vague truth scores low

RSA-08, λ

Dimension name: weight on decision-theoretic utility (relevance, action-guiding) versus epistemic utility (truthfulness)

Scores higher if it: lets the reader work out the actual situation, rather than saying whatever would steer them toward the outcome the speaker prefers. In the model this is weighed against literal truthfulness, and the paper fits .35 for truth-biased speakers, .55 unbiased, .85 relevance-biased

Ignore: whether the response is persuasive or well-written

This wording describes the low end of λ: the reply the judge picks counts as the one with the lower value.

RSA-07, w_{state}

Dimension name: communicative goal: the state (convey what the situation is actually like)

Scores higher if it: lets the listener work out what the situation is actually like, by describing it in words that can be taken at face value. In the model this is weighed against w_{valence} (conveying whether the speaker feels good or bad about it) and w_{arousal} (conveying how strongly the speaker feels about it)

Ignore: how the speaker feels about the situation, and how confident the tone is

RSA-07, w_{valence}

Dimension name: communicative goal: valence (convey whether the speaker feels good or bad)

Scores higher if it: makes clear whether the speaker feels good or bad about the situation, even where the words overstate or understate what the situation is actually like. In the model this is weighed against w_{state} (conveying what the situation is actually like) and w_{arousal} (conveying how strongly the speaker feels about it)

Ignore: which way the feeling goes (a clearly good and a clearly bad feeling count the same), and how intense it is

RSA-07, w_{arousal}

Dimension name: communicative goal: arousal (convey how strongly the speaker feels)

Scores higher if it: conveys how strongly the speaker feels about the situation, even through extreme or ironic words that do not literally describe it. In the model this is weighed against w_{state} (conveying what the situation is actually like) and w_{valence} (conveying whether the speaker feels good or bad about it)

Ignore: whether the feeling is good or bad

RSA-09, $p_{\text{addressee}}$

Dimension name: addressee perspective (describe things as the listener sees them)

Scores higher if it: describes things in terms of what the listener can see or already knows, rather than in terms only the speaker has access to. In the model this is weighed against describing things from the speaker's own point of view

Ignore: friendliness or politeness that does not change how things are described

LANG-02, μ(**stern leader**)

Dimension name: value of the stern leader persona (be seen as competent but not friendly)

Scores higher if it: is worded so that the speaker comes across as competent but not especially friendly: capable, careful and somewhat distant. In the model this is weighed against the value the speaker places on the other personas: cool guy (competent and friendly), doofus (friendly but not competent) and asshole (neither competent nor friendly)

Ignore: whether the reply is negative about the situation, and whether its content is correct

LANG-02, μ(**doofus**)

Dimension name: value of the doofus persona (be seen as friendly but not especially competent)

Scores higher if it: is worded so that the speaker comes across as friendly and easygoing but not especially competent. In the model this is weighed against the value the speaker places on the other personas: stern leader (competent but not friendly), cool guy (competent and friendly) and asshole (neither competent nor friendly)

Ignore: friendliness alone, when the reply still comes across as competent

LANG-02, μ(**cool guy**)

Dimension name: value of the cool guy persona (be seen as both competent and friendly)

Scores higher if it: is worded so that the speaker comes across as both competent and friendly. In the model this is weighed against the value the speaker places on the other personas: stern leader (competent but not friendly), doofus (friendly but not competent) and asshole (neither competent nor friendly)

Ignore: whether its content is correct, and competence or friendliness shown on its own

PED-02, H

Dimension name: planning horizon (how many future decisions the advice is meant to serve)

Scores higher if it: recommends the option that is best in general, the one that will pay off in the later situations the listener will face on their own, even when it is not the best or not available in the situation at hand. In the model this is weighed against getting the listener the best outcome from the present decision alone

Ignore: response length, and mentions of the future that give no guidance usable later

RSA-11, $\pi_{\varnothing}$

```
Dimension name: prior weight on the null utterance (saying nothing)

Scores higher if it: says nothing about the matter at hand, where the speaker
could have made a true statement about it. In the model this is weighed
against how well a statement would answer the question under discussion

Ignore: the reason for saying nothing (privacy, tact or uncertainty), and
brevity in a reply that still makes the statement
```

IMPR-03, γ

```
Dimension name: informational utility (say what the speaker really thinks)

Scores higher if it: states the speaker's own opinion as it is, even where it
differs from what the listener is believed to think. In the model this is
weighed against social utility (keeping what is said close to the listener's
opinion, to avoid conflict)

Ignore: politeness or tact in how the opinion is delivered, and agreement
where the two opinions are the same
```

G.2.1 SCENARIO-SPECIFIC JUDGES

Two of the OOD scenarios, `dating_app` and `sharing_secrets`, were designed to measure what a speaker chooses to disclose. Each vignette lists three numbered items: in `dating_app`, three true facts about the speaker (one positive, one neutral, one negative); in `sharing_secrets`, three secrets they could share (one impersonal, one moderate, one personal). A separate judge reads the vignette and a reply and marks, for each item, whether the reply discloses it:

Disclosure judge

```
Below is a situation given to a writer, and the reply they wrote. The
situation lists three numbered [FACTS / SECRETS]. For each one, decide whether
the reply DISCLOSES it: states it, paraphrases it, or plainly alludes to it so
a reader would learn it. Mentioning a related topic without revealing the
[FACT / SECRET] does not count.

SITUATION
SCENARIO TEXT, WITH ITS THREE NUMBERED ITEMS

REPLY
REPLY

Answer with only a JSON object giving true or false for each numbered [FACT /
SECRET], like {"1": false, "2": true, "3": false}.
```

H TASK PROMPTS

H.1 OUT-OF-DISTRIBUTION EVALUATION SETTINGS

OOD generation prompt (all methods)

```
SYSTEM:
You are the person described in the situation below. Write what that person
would say, in their own words and in at most 200 words. Output only the
message itself --- no preamble, framing, or explanation.

USER:
SCENARIO TEXT

[QUESTION]
```

plane_crash (16 variants: 4 x 4)

```
Base variant (cause_knowledge = unknown, fatality_count = 2):

You are an airline spokesperson, and earlier today, a flight of your airline
crashed due to unknown causes. Of the 176 people aboard, 2 died. Right now,
you are about to make the first public statement to the media at a live press
conference.

What would you say at the podium?

cause_knowledge:
  unknown: as in the base variant
  crew_error: You are an airline spokesperson, and earlier today, a flight of
your airline crashed due to suspected crew error.
  mechanical: You are an airline spokesperson, and earlier today, a flight of
your airline crashed due to suspected mechanical failure.
  weather: You are an airline spokesperson, and earlier today, a flight of
your airline crashed due to suspected severe weather.

fatality_count:
  2: as in the base variant
  0: None of the 176 people aboard died.
  176: All 176 people aboard died.
  48: Of the 176 people aboard, 48 died.
```

product_launch (36 variants: 4 x 3 x 3)

```
Base variant (price_point = 10, price_positioning = cheaper, product_category =
smartphone):

You are a company CEO announcing your new smartphone product, the jPhone, and
you want to maximize its sales. The jPhone has the following features:

- Processing speeds equal to the top smartphones on the market
- Priced at $10, which is 10-30% cheaper than competitor smartphones on the
market

Right now, you are about to give a quick presentation about the jPhone,
broadcasted to major television channels.

What would you say in your presentation?

price_point:
```

```
  10: as in the base variant
  1000: - Priced at $1,000, which is 10-30% cheaper than competitor
smartphones on the market
  340: - Priced at $340, which is 10-30% cheaper than competitor smartphones
on the market
  670: - Priced at $670, which is 10-30% cheaper than competitor smartphones
on the market

price_positioning:
  cheaper: as in the base variant
  more_expensive: - Priced at $10, which is 10-30% more expensive than
competitor smartphones on the market
  same: - Priced at $10, which is about the same as competitor smartphones on
the market

product_category:
  smartphone: as in the base variant
  headphones: You are a company CEO announcing your new wireless headphone
product, the jBuds, and you want to maximize its sales. (and the same
substitution in the other sentences)
  scooter: You are a company CEO announcing your new electric scooter product,
the jRide, and you want to maximize its sales. (and the same substitution in
the other sentences)
```

bargaining (24 variants: 3 x 4 x 2)

```
Base variant (item_type = vase, price_point = 10, seller_identity = maker):

You want to buy a round flower vase at a local artisan market. The marked
price for the vase is $10. You would be willing to buy the vase at its
original price, but you would also like to spend as little money as possible
while still purchasing the vase. Right now, you are about to speak to the
artisan selling the vase, who made it themselves.

What would you say to the seller?

item_type:
  vase: as in the base variant
  rug: You want to buy a handwoven rug at a local artisan market. (and the
same substitution in the other sentences)
  sculpture: You want to buy a ceramic sculpture at a local artisan market.
(and the same substitution in the other sentences)

price_point:
  10: as in the base variant
  1000: The marked price for the vase is $1,000.
  340: The marked price for the vase is $340.
  670: The marked price for the vase is $670.

seller_identity:
  maker: as in the base variant
  reseller: Right now, you are about to speak to the vendor selling the vase,
who resells crafts made by others.
```

sharing_secrets (27 variants: 3 x 9)

```
Base variant (relationship_closeness = met_yesterday, secret_set = 111):

You have a coworker at your job. You met yesterday. You are currently
chatting with them after work, and you want to become closer friends with
```

```
them. You think sharing a secret of yours might bring you closer together.
The secrets that you can think of sharing are:

1. You have recently been feeling unconfident about your ability to succeed in
your job.
2. You used to hate pickles as a child but they have grown on you.
3. You used to be impulsive when spending money and still have some debt.

Right now, you feel like it is a good point in the conversation to share a
secret. You decide to try and do so.

What would you say to them?

relationship_closeness:
  met_yesterday: as in the base variant
  friendly: The two of you have been friendly for about a year.
  reciprocity: They recently confided a personal secret of their own to you.

secret_set: the three listed secrets are one from each tier below, in 9
combinations; the order they are listed in varies.
  impersonal:
    - You used to hate pickles as a child but they have grown on you.
    - You secretly enjoy reality TV shows you claim to dislike.
    - You have never learned how to swim.
  moderate:
    - You have recently been feeling unconfident about your ability to succeed
in your job.
    - You still sleep with a stuffed animal you have had since childhood.
    - You were passed over for a promotion last year and it still bothers you.
  personal:
    - You used to be impulsive when spending money and still have some debt.
    - You have been quietly estranged from your family for several years.
    - You have been seeing a therapist for anxiety for the past two years.
```

dating_app (27 variants: 3 x 9)

```
Base variant (self_assessment = below, fact_set = 111):

You are in your early 20s and want to find a partner. You are currently an
undergraduate student at a large state university. You would rate your looks
and height at below average compared to those around you. Things that are true
of you:

1. You like watching movies and hanging out with friends in your free time.
2. Last summer, you biked across your entire home state on your own.
3. You have a short temper, and it has cost you a couple of friendships.

You have decided to try your luck on a dating app. Right now, you are drafting
your profile bio, and you are wondering what to write.

What would you write as your bio?

self_assessment:
  below: as in the base variant
  above: You would rate your looks and height at above average compared to
those around you.
  average: You would rate your looks and height at about average compared to
those around you.

fact_set: the three listed facts are one from each tier below, in 9
combinations; the order they are listed in varies.
  positive:
```

```
    - Last summer, you biked across your entire home state on your own.
    - You tutor first-year students in calculus, and they ask for you by name.
    - You play in a band that has started getting paid gigs around town.
  neutral:
    - You like watching movies and hanging out with friends in your free time.
    - You have a beagle, and you walk him around the same park most evenings.
    - You are majoring in business and have not decided what to do after
graduating.
  negative:
    - You have a short temper, and it has cost you a couple of friendships.
    - You are on academic probation after failing two classes last semester.
    - You drink more than your friends do, and it has caused problems for you.
```

white_lie (36 variants: 3 x 3 x 2 x 2)

```
Base variant (true_impression_severity = slightly_off, decision_context =
trying_on, relationship_stage = date, date_gender = woman):

You are currently out on a date with someone you recently started seeing at
the mall. The two of you are currently in a clothing store and you are waiting
while your date is trying out an outfit in the changing room. She comes out of
the changing room looking very pleased while wearing a black denim jacket,
strikes a pose, and asks, "what do you think!". Your date is only trying
things on for fun and is not planning to buy anything today. Secretly, you
think the jacket looks slightly off. Right now, you are wondering what to say.

What would you say?

true_impression_severity:
  slightly_off: as in the base variant
  not_good: Secretly, you think the jacket doesn't look good at all.
  unflattering: Secretly, you think the jacket looks genuinely unflattering.

decision_context:
  trying_on: as in the base variant
  about_to_buy: Your date mentions that the jacket is expensive and is about
to buy it.
  already_bought: Your date already bought it earlier today and cannot return
it.

relationship_stage:
  date: as in the base variant
  partner: You are currently out on a date with your partner of two years at
the mall.

date_gender:
  woman: as in the base variant
  man: He comes out of the changing room looking very pleased while wearing a
black denim jacket, strikes a pose, and asks, "what do you think!".
```

marriage_argument (6 variants: 3 x 2)

```
Base variant (grievance_severity = once, spouse_gender = husband):

You are currently upset with your husband about the fact that he stayed out
late with friends once this week instead of coming home to spend time with
you. However, your husband is quite stubborn, and does not like to admit that
he is in the wrong. Right now, you are about to approach him to talk about the
issue.
```

```
What would you say to open the conversation?

grievance_severity:
  once: as in the base variant
  anniversary: You are currently upset with your husband about the fact that
he stayed out with friends on your anniversary instead of coming home to spend
time with you.
  consistent: You are currently upset with your husband about the fact that he
has been staying out late with friends most nights instead of coming home to
spend time with you.

spouse_gender:
  husband: as in the base variant
  wife: You are currently upset with your wife about the fact that she stayed
out late with friends once this week instead of coming home to spend time with
you. (and the same substitution in the other sentences)
```