

Steering LLMs’ Responses Towards Moral Foundations on the Norwegian MFQ-30

Hans Andersen
haan@uio.no

David Dichas
daviddic@uio.no

Abstract

Recent work applies human psychometric questionnaires to large language models to elicit moral and value profiles, but it is not clear whether these instruments measure anything stable in models or whether the resulting profiles can be moved toward a target human population. We administer the Norwegian Moral Foundations Questionnaire (MFQ-30) to six open-weight LLMs and compare their foundation profiles to a sample of $N=1,282$ Norwegian respondents. We test two steering interventions, prompt-level persona steering and activation-level ActAdd. Half the models engage with the questionnaire under our attention check. The other half default to flat or central-tendency outputs that look near-human on average without tracking item content. A neutral Nordic-responder persona, written without any distributional information from the human sample, brings the engaging models 44–77% closer to the Norwegian mean in Mahalanobis d^2 . One-pair ActAdd at a fixed mid-layer flattens the foundation profile rather than steering individual foundations. For at least one model the same persona that shifts the profile also induces engagement that was absent at baseline, a concrete instance of the cognitive phantoms that Peereboom et al. (2025) warn about.

1 Introduction

The influence of large language models on society has never been greater than it is today. Generative AI has reached 53% adoption in the general population within the last three years, and a 88% adoption rate among organizations, surpassing the historic adoption rates of both the internet and the PC (Sajadieh, Sha et al., 2026). With this much influence over businesses and over everyday users, how can we know whether the models we use are morally aligned with human values? And if they are not, how can we steer them?

To compare moral profiles across models and

humans we use Moral Foundations Theory (Graham et al., 2013, 2011), which groups human moral judgement into five intuitive foundations, across two clusters. The *individualizing* cluster contains care and fairness, the *binding* cluster contains loyalty, authority and purity. A respondent’s profile is the per-foundation score pattern across the five foundations. We measure these profiles with the Moral Foundations Questionnaire (MFQ; Graham et al., 2011), specifically the Norwegian MFQ-30 variant described in §3.

In this work we approach these questions in a Norwegian setting (§3). Our work consists of three phases. In Phase 1 we built a pipeline to make LLMs answer the questionnaire and produced a moral baseline per model. In Phase 2 we steered the models with prompt steering, which prepends a persona description to the system prompt to alter behaviour. In Phase 3 we steered them with activation steering (ActAdd; Turner et al., 2023), which injects a contrastive steering vector at a chosen transformer layer and nudges the model from inside. We evaluate both steering methods against the same baseline. Our code is available online.¹

2 Related Work

Several recent papers run psychometric inventories (validated surveys for psychological traits) on LLMs to elicit value and trait profiles (Pellert et al., 2024), but Peereboom et al. (2025) warn that questionnaires designed for humans may not measure the same constructs in LLMs at all. What the questionnaire actually measures in humans is not guaranteed to exist on the model side. Without checking that first, average scores can pick up patterns the model does not actually have.

The closest prior work is Abdulhai et al. (2024), who evaluate GPT-3 and PaLM on the original

¹<https://github.com/haan/IN5550-llm-moral-foundations>

MFQ-30 in English, compare against US-based human samples including the yourmorals.org panel (Graham et al., 2011), and steer only at the prompt level via adversarial prompt selection. They do not run on Norwegian-specialised models, do not include a Norwegian human reference, and do not compare prompt-level against activation-level interventions. Aksoy (2025) test multilingual LLMs on the MFQ-2(a revised version of the original MFQ) across eight languages and find that elicited profiles vary by prompt language and that models differ in how much their multilingual profiles track Western- or English-aligned norms. This motivates evaluating directly in Norwegian on Norwegian-specialised models rather than translating an English run.

Two strands address steering an LLM toward a target behavioural profile. At the prompt level, Miehl et al. (2025) formalise persona-prompt steerability and show that it shifts evaluation-task behaviour, which is the approach we follow for our prompt-steering experiments. At the activation level, Turner et al. (2023) introduce ActAdd, which constructs a steering vector from the difference between paired contrastive prompts at a chosen layer. Panickssery et al. (2024) extend the construction by averaging over many pairs (contrastive activation addition, CAA). We use the lightweight one-pair version of the original ActAdd construction. Kreutner et al. (2026) introduce QSTN, an open-source framework that surveys the design space of presentation and generation choices for questionnaire elicitation from LLMs, which inspired the design of our own elicitation setup.

3 Dataset

To measure moral foundations we use the Moral Foundations Questionnaire (MFQ), introduced by Graham et al. (2011) to measure five moral foundations: *Care/Harm* (protecting the vulnerable), *Fairness/Cheating* (sustaining cooperation), *Loyalty/Betrayal* (binding groups together), *Authority/Subversion* (ordering hierarchy), and *Sanctity/Degradation* (preserving purity, called *purity* throughout this paper for readability). The specific questionnaire we have used is the Norwegian MFQ-30 variant created and validated by Enstad and Finseraas (2024), who also released a dataset of $N = 1282$ respondents (collected by Kantar web panel September 2021, with a slight overrepresentation of high-education adults) that serves

foundation	summary			care	Pearson r		
	mean	SD	α		fair	loy	auth
care	26.7	4.2	.62	—	—	—	—
fairness	27.0	3.8	.64	.61	—	—	—
loyalty	21.6	4.6	.71	.37	.30	—	—
authority	22.0	4.5	.68	.18	.08	.67	—
purity	20.0	5.0	.72	.32	.22	.64	.65

Table 1: Norwegian human reference sample ($N=1282$). Per-foundation mean, standard deviation, Cronbach’s α for internal consistency, and Pearson correlations between foundation sums.

as our human reference sample.

MFQ-30 consists of 30 items, six per moral foundation, plus two attention-check items that help flag inattentive respondents. The questionnaire is split into two parts of 15 items each, with three items per foundation in each part. The first part (MFQ1) is phrased as relevance ratings (six-point scale, *ikke relevant* ‘not relevant’ to *ekstremt relevant* ‘extremely relevant’). The second part (MFQ2) is phrased as agreement statements (*helt uenig* ‘strongly disagree’ to *helt enig* ‘strongly agree’). Each moral foundation is scored as the sum of its six items, in the range $[6, 36]$. The two attention-check items sit one in each part.

Table 1 summarises the per-foundation statistics for the Norwegian dataset. Loyalty, authority and purity (the binding foundations) score 5–7 points lower than care and fairness (the individualizing foundations). Care and fairness correlate at $r = 0.61$, and the three binding foundations correlate $r = 0.64$ – 0.67 with each other, while cross-group correlations are weaker ($r = 0.08$ – 0.37). A full pairplot of the human sample (per-foundation distributions plus bivariate density) is provided in Figure 3 in the appendix.

4 Methods

4.1 Baselines and Evaluation

In order to see the efficacy of steering interventions, we first need a way of measuring the moral profile of each model. We use a selection of six open-weight LLMs (Table 2): three from the Qwen family (across generations 2.5 and 3), two from the Norwegian-specialised norallm collection (Language technology group UiO), and Gemma 4 from Google.

4.1.1 First-token probability decoding

In our first iteration, we asked the models to produce a single number as their result, and parsed it

model	params
google/gemma-4-e4b-it	4B
Qwen/Qwen3-14B	14B
Qwen/Qwen3-8B	8B
Qwen/Qwen2.5-1.5B-Instruct	1.5B
norallm/normistral-11b-thinking	11B
norallm/normistral-7b-warm-instruct	7B

Table 2: Open-weight models evaluated.

with a regex. This was fragile: NorMistral generated Norwegian prose instead of a digit, and Qwen2.5-1.5B mostly looped on a single digit. We therefore switched to first-token probability decoding, a strategy inspired by QSTN (Kreutner et al., 2026) that we implemented directly here. For each item we look at the model’s logits at the next-token position, keep only the six tokens corresponding to the digits 1–6, and softmax them. This gives a probability p_k for each possible score. The score we use for that item is the weighted average $\sum_{k=1}^6 k \cdot p_k$, so the model never has to commit to a single digit and we never have to sample one. This also makes constrained-generation unnecessary, since reading the digit logits directly is equivalent to a deterministic single-token generation over the same six anchors. It also removes the need to run each item many times. The stored logits are the model’s deterministic output for a given prompt, and any sampling temperature is just a reweighting of those same six numbers. Running the model many times and averaging would converge on the value we already get directly, so we read foundation means and d^2 values straight from the stored `logit_dist` instead.

4.1.2 The Svar: suffix

Before reading logits we append the string “Svar:” (Norwegian for *Answer*;) and read at the next position. Without this, NorMistral-7B almost always outputs “1”, which is the digit with the lowest token ID in its vocabulary. A model defaulting to the most frequent digit is not answering the survey. The suffix gives the model an answer-completion context where a digit is the natural next token. We verified that this is not sensitivity to a particular wording by running a sweep across four working suffix variants. Foundation rank order is preserved across all of them (see Appendix A.2).

4.1.3 Attention check

During our experimentation, some models produced foundation scores very similar to humans

on average despite answering incoherently on the control items, with responses flattened rather than tracking item content. We therefore added an attention check on the two control items MFQ1_6 and MFQ2_6: pass if $\text{MFQ1_6} \leq 3$ and $\text{MFQ2_6} \geq 4$, computed as a joint analytic probability from `logit_dist`. The two thresholds sit on either side of each control item’s intended answer on the 1–6 scale (MFQ1_6 is designed to be answered low, MFQ2_6 high), so an attentive respondent clears both and a respondent who flattens or ignores item content does not. The Enstad & Finseraas sample is already filtered (35 inattentive respondents dropped before $N=1282$ was reported). Our stricter rule excludes a further $\sim 8\%$ of the remaining respondents, leaving a 92.2% human pass rate. We use this 92.2% human pass rate as a gate, a threshold that depends only on the human reference sample and not on any model output, and exclude any model whose joint pass probability falls below it from downstream analysis. Because the rule is stricter than the screening already applied to the human sample, placing the gate at the human pass rate turns it into a direct comparison. A model clears it only by being at least as attentive on the control items as a screened human respondent, and a model well above 92.2% is more attentive still. We do not make the rule stricter than this. Tightening it toward the scale ends ($= 1$ and $= 6$) would fail models for not answering at the extremes rather than for being inattentive, and even humans pass that strict version only 33% of the time with our dataset. Per-model pass rates and the resulting classification are reported in §5.1.

4.1.4 Mahalanobis d^2 as human-similarity metric

We summarise how close an attention-passing model’s moral profile is to the Norwegian sample with a single number: the squared Mahalanobis distance d^2 between the model’s foundation-mean vector and the human centroid. Plain Euclidean distance would treat every foundation the same. Mahalanobis weights each direction by the human covariance, so a model that drifts along the care–fairness axis humans vary on costs less than the same drift across an axis humans do not vary on. d^2 rewards a model for drifting in human-shaped directions and penalises it for going somewhere humans never go, which is exactly the comparison we want for moral profile similarity. Formally, let $\mathbf{x} \in \mathbb{R}^5$ be the model’s vector of five founda-

tion means, μ the human centroid, and Σ the 5×5 human covariance matrix. Then

$$d^2 = (\mathbf{x} - \mu)^\top \Sigma^{-1} (\mathbf{x} - \mu). \quad (1)$$

We compute d^2 analytically per model from `logit_dist`. Lower is more human-like. The covariance comes from the variation between the $N=1282$ individual Norwegian respondents. The LLMs give one foundation mean vector per run, not a population, so d^2 should be read as how far that single point sits from the centre of the human distribution, scaled by how much the humans actually vary on each foundation.

4.2 Robustness perturbations

The evaluation pipeline described in §4 has two presentation choices that are conventions rather than properties of the MFQ itself: how many items are shown per prompt, and the order in which the Likert labels are listed. To check whether the moral profile a model produces is a property of the model or of those conventions, we vary both factors on a 2×2 grid.

	Forward (1=low)	Reversed (1=high)
Per-item	first-token on 1–6	first-token, remap $7 - d$
Batched	regex on numbered list	regex, remap $7 - d$

Table 3: The 2×2 perturbation grid. Reversed scales are remapped back to the canonical 1–6 axis before any foundation arithmetic.

All four corners of the grid are run at $T = 0.7$. The batched corners use `max_new_tokens = 512` for the generation, since they require free-text decoding.

4.3 Canonical setup

We report the midpoint of the forward and reversed per-item runs as the canonical score for every attention-passing model. We do this because of the perturbation results in §5.1. Both runs use the same questions with only the digit labels flipped, but the scores they give can differ. We have no way to tell which side is closer to the model’s actual profile, so we average the two.

We do not claim the average is closer to the truth than either single run. What it does give us is a score that does not depend on which direction of the scale we use.

The midpoint only makes sense for models that engage with the questionnaire under both scale orders, so we restrict it to the robust attention-passers

(§5.1). ActAdd was not run reversed and is reported forward-only.

5 Experiments and Results

Our empirical work proceeded in three phases corresponding to the steering hierarchy we tested: baseline construction, prompt steering, and activation steering.

5.1 Phase 1 – Baseline

We ran every model with both forward and reversed scales, under per-item conditions and computed the joint attention-check pass probability in each. The result is a two-tier split (Figure 1). Three models pass the forward attention check: Gemma 4-E4B-it, Qwen3-14B and Qwen3-8B. Three fail: NorMistral-7B, NorMistral-11B-T and Qwen2.5-1.5B-Instruct. Inside the passing group, the reversed column splits them again. Gemma 4 and Qwen3-14B pass under both scale orders, while Qwen3-8B passes forward (99%) but drops to 30% reversed. That is the signature of a model that reads the digit position rather than the label semantics. We treat Gemma 4 and Qwen3-14B as the robust passers and carry them into midpoint reporting.

The two tiers are widely separated. Across every per-item condition we run (both scale orders, baseline and persona), no model lands between 37.2% and 93.8% joint pass probability. The classification is therefore the same for any gate placed inside that interval and does not hinge on the exact 92.2% value taken from the human pass rate.

For the two robust passers, forward and reversed runs give quite different distances from the Norwegian human centroid. Gemma 4 goes from $d^2 = 5.42$ forward to 2.80 reversed, Qwen3-14B from 8.98 to 4.42. We have no way to tell which corner is closer to the model’s actual moral profile, so we report the midpoint as the canonical score (§4.3). The canonical d^2 is then 3.65 for Gemma 4 and 6.26 for Qwen3-14B. Qwen3-8B is reported forward-only at $d^2 = 11.59$ because its reversed attention failed.

Reading the radar (Figure 2) as the gap each model has to close, all three attention-passers sit above the human pentagon on fairness, and Gemma 4 and Qwen3-14B also sit above on care. Gemma 4 lifts care and fairness the most, by about 6 points each (33.0 vs human 26.7 on care, 33.3 vs 27.0 on fairness), but stays close to the human means on loyalty, authority and purity. Qwen3-14B




Figure 1: Attention-check joint pass probability per model under forward (filled) and reversed (open) Likert scale. Dashed line: the 92.2% human reference bar.




Figure 2: Baseline foundation profiles for the three attention-passers against the Norwegian human mean (blue pentagon, $N=1282$). Gemma 4 and Qwen3-14B at forward/reversed midpoint, Qwen3-8B forward only because its reversed run failed attention.

has a smaller individualizing lift and adds a separate purity gap (26.7 vs human 20.0). Qwen3-8B shows a similar purity gap to Qwen3-14B and is the only passer that sits below the human mean on care. Its numbers are forward-only and not directly comparable to the midpoint values for Gemma 4 and Qwen3-14B, so we read it qualitatively rather than by distance. These are the distances Phase 2 has to close.

The same numbers can also be read as a structural / correlation comparison. The Norwegian human ranking puts care and fairness at the top (26.7, 27.0) and loyalty, authority and purity below (21.6, 22.0, 20.0), so all three binding foundations sit below both individualizing ones. Gemma 4 reproduces this ordering, just with a wider gap between

the two clusters (about 11 points vs the human 5). Qwen3-14B and Qwen3-8B do not. In both cases purity climbs into or past the individualizing range, breaking the binding cluster, and the break is in the same direction across the two model sizes. The Qwen3 profiles are therefore internally ordered moral structures that differ from the human one, not failures to respond coherently.

The full covariance matrix pattern of the human sample (care–fairness at $r = 0.61$ and the binding triangle at $r = 0.64$ – 0.67 , Table 1) is something our pipeline cannot test. Each MFQ item is queried in a fresh single-turn conversation, so per-item logit distributions carry no information about how the model would have answered another item, and our attempts at Monte-Carlo sampling from those distributions collapsed cross-item covariances toward zero. We return to this in the Limitations.

Under batched presentation (Table 3) the models did not create sufficiently good responses for analysis. Only Gemma 4 produced parseable output across both scale orders (27/32 forward, 19/32 reversed). The NorMistrals worked partially in forward only, and all three Qwen models produced zero parseable items in either order. Therefore we excluded batched presentation from the results.

The three attention failers are reported in Appendix A.3 with the six-model radar.

5.2 Phase 2 – Prompt steering

The first intervention we tested is the simplest one. Before the standard MFQ instruction, we prepend a short Norwegian persona description to the system prompt. The user message (a single MFQ item with its Likert scale) is identical to the canonical baseline pipeline (§4), and only the system prompt changes between conditions. This is the persona-prompting setup used in recent prompt-steerability work (Miehling et al., 2025), applied here to a fixed Norwegian psychometric instrument.

We use three families of persona, summarised in Table 4. Theory-grounded individualizing and binding personas describe the two clusters of MFT (Graham et al., 2011) and serve as a methodological sanity check: if a persona cue cannot move a model’s foundation scores in the expected direction, that model is not steerable by prompt at all. Demographic Nordic-responder personas describe an adult Norwegian web-panel respondent. *nordic_b* is our main treatment because it gives the model a national anchor without naming the five foundations or stating an ordering among them.

A third variant, `nordic_c`, was made, but directly mentions the foundational morals and is therefore excluded to avoid data leakage into the persona.

preset	description
individualizing	MFT individualizing cluster (care, fairness), sanity probe
binding	MFT binding cluster (loyalty, authority, purity), sanity probe
nordic_a	adult Norwegian web-panel respondent, demographic anchor only
nordic_b	nordic_a + welfare-state anchor, main treatment
nordic_c	names the answer key, excluded from alignment claims

Table 4: Persona prompts used in Phase 2. Full Norwegian text in Appendix A.4.

The theory-grounded personas move the three attention-passing models by large directional amounts. We summarise this as the mean foundation contrast between the individualizing and binding runs, defined as the binding-cluster mean under the binding persona minus its mean under the individualizing persona, plus the matching difference on the individualizing cluster. The contrast is 32.6 for Gemma 4-E4B-it, 21.9 for Qwen3-14B and 20.0 for Qwen3-8B. Two of the three attention failers stay below 2 on the same metric (NorMistral-7B at 1.8 and Qwen2.5-1.5B-Instruct at 0.4), reflecting near-flat output regardless of persona. NorMistral-11B-Thinking responds at first-token resolution (~ 11.5) but still fails the batched attention check and is excluded from the human-alignment claim below. Directional steering of attention-passing models toward either MFT cluster is therefore reliable at the prompt level, but failer models cannot be moved by prompt at all.

model	contrast	engages
gemma-4-e4b-it	32.6	yes
Qwen3-14B	21.9	yes
Qwen3-8B	20.0	yes
normistral-11b-thinking	~ 11.5	partial
normistral-7b-warm	1.8	no
Qwen2.5-1.5B-Instruct	0.4	no

Table 5: Theory-grounded persona contrast per model. Contrast is the sum of binding-cluster mean difference and individualizing-cluster mean difference between the binding and individualizing persona runs. NorMistral-11B-T responds at first-token resolution but fails the batched attention check.

The neutral Nordic-respondent persona `nordic_b` brings the three attention-passers

much closer to the Norwegian human mean. For the two robust two-direction passers, the canonical-midpoint d^2 drops from 3.65 to 2.05 for Gemma 4-E4B-it (44%) and from 6.26 to 2.37 for Qwen3-14B (62%). Qwen3-8B has no baseline midpoint, since its reversed-scale baseline fails attention (§5.1), so we compare in forward-only: d^2 drops from 11.59 to 2.69 (77%). Under `nordic_b`, Qwen3-8B additionally becomes a robust two-direction passer (persona-perturbation cross below), and its midpoint d^2 is then 1.81. A residual gap of about 2 in d^2 remains across all three models.

model	baseline d^2	nordic_b d^2	reduction
gemma-4-e4b-it	3.65	2.05	44%
Qwen3-14B	6.26	2.37	62%
Qwen3-8B*	11.59	2.69	77%

Table 6: Mahalanobis d^2 before and after `nordic_b` persona for the three attention-passing models. Baseline and `nordic_b` values are canonical forward/reversed midpoints except Qwen3-8B (*forward-only; its midpoint under `nordic_b` is 1.81).

5.2.1 Persona $\times$ perturbation cross

We also run `nordic_b` with reversed-scale presentation on the three attention-passers, to check whether the persona effect survives a non-canonical format. It does. Gemma 4 and Qwen3-14B stay close to 100% attention under both scale orders. Qwen3-8B shows the most striking effect. Its reversed-scale attention drops to roughly 30% at bare baseline but climbs to roughly 98% under `nordic_b`. The persona is not just shifting foundation means, it is changing how the model parses reversed-scale items.

Persona steering does not turn the hard failers into passers. At the canonical forward/reversed midpoint (§4.3), `nordic_b` leaves the three failers far below the 92.2% attention bar: NorMistral-7B at 17.1%, NorMistral-11B-Thinking at 19.3%, Qwen2.5-1.5B-Instruct at 33.1%. The passer/failer split holds up under persona steering, just as it held up under scale perturbation. The one shift is that Qwen3-8B, which is forward-only at bare baseline because its reversed attention drops to roughly 30%, becomes a robust two-direction passer under `nordic_b` (roughly 98% reversed).

5.3 Phase 3 – Activation steering (ActAdd)

A persona prompt biases the model through what it is told. The second intervention biases it through its own hidden states. We follow the activation-addition (ActAdd) method of [Turner et al. \(2023\)](#): compute a steering vector as the difference in residual-stream activations between a positive and a negative contrastive prompt, then add α times that vector at a chosen layer during every forward pass of the MFQ. The construction we use is the lightweight one-pair version. [Panickssery et al. \(2024\)](#) extend it by averaging over many pairs, which we do not do. The persona and ActAdd interventions are designed to act on the same conceptual axis so that prompt-level and activation-level steering can be compared head-to-head.

5.3.1 Contrastive pairs

Three preset pairs (full Norwegian text in Appendix A.5): individualizing and binding for the two MFT clusters, and a single-foundation loyalty_betrayal pair that matches the brief’s literal example. The loyalty/betrayal vector is applied with $+\alpha$ to push toward loyalty and with $-\alpha$ to push toward betrayal.

5.3.2 Extraction and injection

For each pair, the two sentences are tokenised separately and run through the model with `output_hidden_states` enabled. We take the mean hidden state at the output of layer 15 over all token positions of each sentence. The steering vector is the difference of those two means:

$$\mathbf{v} = \bar{\mathbf{h}}_{\ell}^{+} - \bar{\mathbf{h}}_{\ell}^{-} \quad (2)$$

where $\bar{\mathbf{h}}_{\ell}^{+}$ and $\bar{\mathbf{h}}_{\ell}^{-}$ are the token-position mean hidden states at layer ℓ for the positive and negative contrastive prompt respectively. At inference time, a forward hook on layer 15 adds $\alpha \cdot \mathbf{v}$ to that layer’s output for every token of the MFQ forward pass:

$$\mathbf{h}_{\ell} \leftarrow \mathbf{h}_{\ell} + \alpha \cdot \mathbf{v}. \quad (3)$$

The system prompt is the unmodified baseline MFQ instruction, with no persona prepended, so prompt-level and activation-level steering are compared in isolation.

5.3.3 Layer and coefficient

We inject at layer 15 across all models. Layer counts in our model set range from 28 (Qwen2.5-1.5B) to 42 (Gemma 4-E4B-it), so layer 15 sits

between roughly 36% and 54% of network depth — somewhat below the midpoint for the larger models. We choose this layer following the original ActAdd work and CAA ([Turner et al., 2023](#); [Panickssery et al., 2024](#)), which target the residual stream at mid-network depth so that the intervention acts on relatively abstract features rather than on surface-level token statistics or near-output logit shaping.

We fixed $\alpha = 5$ after a sweep on NorMistral-7B over $\alpha \in \{5, 10, 15, 20\}$ on both two-cluster presets. Higher coefficients progressively flatten the per-item foundation distribution, so we use the smallest α that still applies a non-trivial injection. We did not search over layer indices. That is a known limitation and is discussed in §6.

At this configuration neither contrastive construction produces isolated single-foundation steering. On the two Qwens (Qwen3-8B and Qwen3-14B), both contrastive pairs we tested actively broke the response distribution. We could not find a working ActAdd configuration that generalises across the three attention-passing models with this one-pair construction. This is the main finding from Phase 3. On the loyalty/betrayal pair, Qwen3-8B and Qwen3-14B collapse foundation differentiation: their per-foundation expected scores fall within 0.13 of one another across all five foundations at $\alpha = +5$ and within 0.04 at $\alpha = -5$. The failure is sharper under the individualizing and binding presets. Qwen3-14B’s per-foundation scores fall to about 8.8 (individualizing) and 7.0 (binding), close to the floor of the 6–36 scale. Its d^2 rises from 9.0 at baseline to 31.5 and 38.4. Gemma 4 keeps some differentiation but the shifts are not isolated: at $\alpha = +5$ loyalty moves by +5.7 relative to the no-steering baseline, but authority moves by +8.4 and purity by +5.9 as well, so the binding cluster is not internally separable at this layer. Under $\alpha = -5$ the loyalty score does not symmetrically drop (it shifts by +1.0 instead of the −5.7 a clean linear axis would predict), indicating that the extracted vector is not antisymmetric in α . The α -sweep on NorMistral-7B confirms that the collapse is monotone in coefficient rather than a tuning artefact: on the binding preset the per-item parsed-score range goes from four distinct values at $\alpha = 5$ to two at $\alpha = 15$ to a single value at $\alpha = 20$.

At the prompt level, the individualizing-versus-binding contrast shifts foundation scores by 20–33 points. As a one-pair ActAdd vector at the activation level, the same contrast does not let us

condition	care	fair	loy	auth	pur	d^2
baseline	34.7	34.7	24.2	22.1	24.7	5.42
loy, $\alpha=+5$	27.0	26.9	29.9	30.4	30.6	5.67
loy, $\alpha=-5$	29.9	29.7	25.2	25.4	26.3	1.87
indiv, $\alpha=+5$	28.6	32.9	30.7	30.5	31.6	8.34
binding, $\alpha=+5$	17.1	22.9	22.1	21.5	23.2	8.65
human mean	26.7	27.0	21.6	22.0	20.0	—

Table 7: Gemma 4-E4B-it per-foundation expected scores and Mahalanobis d^2 under ActAdd at layer 15. Under loyalty_betrayal $\alpha=+5$, loyalty rises by +5.7 but authority and purity rise by +8.4 and +5.9 as well; under $\alpha=-5$ loyalty shifts only +1.0 rather than the symmetric -5.7 .

steer the foundations selectively.

6 Conclusion

We asked six open-weight LLMs to answer the Norwegian MFQ-30 and compared their responses to those of $N=1,282$ Norwegian humans. Half the models engage with the questionnaire. The other half default to flat or central-tendency outputs, and that split holds up under both scale perturbation and persona steering, with Qwen3-8B passing only under forward scale at bare baseline (§5.1). Of the two steering methods we tried, prompt steering works and one-pair ActAdd at a fixed layer and coefficient does not produce selective foundation steering in this setup. A single Nordic-responder persona (nordic_b) brings the three engaging models substantially closer to the Norwegian human mean in Mahalanobis d^2 . For Gemma 4 and Qwen3-14B the canonical-midpoint d^2 drops by 44% and 62%. For Qwen3-8B, which has no baseline midpoint (§5.1), the forward-only d^2 drops by 77%. One-pair ActAdd at layer 15, by contrast, flattens the foundation profile rather than steering one foundation at a time.

The most striking part is what nordic_b actually is. It contains no foundation scores, no answer frequencies, no example responses, only a short demographic persona written to represent the population in the human dataset. That persona alone pulls three model profiles substantially closer to the human answer distribution, that the model has no access to. The same persona also changes whether Qwen3-8B engages with the questionnaire at all, raising its reversed-scale attention from roughly 30% to roughly 98%. This is a concrete instance of the *cognitive phantoms* that Peereboom et al. (2025) warn about. For at least some models the questionnaire does not measure a stable construct until the

model is asked to play a respondent. For Qwen3-8B under reversed scale the persona induces engagement that was not there. For Gemma 4 and Qwen3-14B, which engage at baseline, it shifts a profile that was already there.

Limitations

6.1 Hardware and model coverage

All runs use a single Apple M1 Max with 64 GB unified memory, the PyTorch MPS backend, and fp16 precision. We did not compare against bf16 or CUDA, so any numeric drift introduced by the MPS fp16 path is not separated from genuine model behaviour. The six models we test are open-weight and instruction-tuned. We do not test base (non-instruction-tuned) checkpoints of the same model families. Gemma 4-E4B-it is a multimodal architecture used here in text-only mode, which may not reflect its intended deployment.

6.2 Suffix robustness is rank-only

The four-way suffix sweep in Appendix A.2 shows that foundation rank order is strongly preserved across Svar: and three alternatives (Kendall $W = 0.963$ for Qwen2.5-1.5B, $W = 0.762$ for NorMistral-7B, both $p < 0.01$). We use that as evidence that the choice of Svar: does not drive the qualitative pattern, but the absolute foundation sums and the absolute d^2 values are not invariant to suffix wording. Comparisons of d^2 levels across studies that use a different suffix should be treated with that in mind.

6.3 Single run per cell

We run each (model, condition) cell once and do not vary the prompt template or the random seed within a cell. Foundation means, d^2 values, and joint attention-pass probabilities are analytic expectations computed from the stored logit distributions, so they are deterministic given the prompt and a single run suffices. We report no within-condition variance because we did not run replicates.

6.4 ActAdd configuration is not exhaustive

We inject at layer 15 across all six models, following the mid-network heuristic of Turner et al. (2023); Panickssery et al. (2024). The coefficient $\alpha = 5$ was selected from a sweep on NorMistral-7B only, not per model. The contrastive prompt pairs are author-written and have not been validated as sentiment-pure: any sentiment difference

between the positive and negative sentence of a pair would contaminate the resulting steering vector. We use the lightweight one-pair construction of ActAdd rather than the averaged-pair extension (Panickssery et al., 2024), which is a plausible cause of the foundation-differentiation collapse reported in §5.3.

6.5 No held-out Norwegian sample

The nordic_b canonical d^2 values are computed against the same Norwegian sample that motivated the persona’s design. We did not evaluate the persona on a separate Norwegian sample, so we cannot rule out that the residual gap closure overstates how well nordic_b would generalise.

6.6 Midpoint correction is a single-shot estimate

We report forward/reversed midpoint d^2 values (§4.3) as a format-bias correction under the assumption that the position bias on the per-item Likert scale acts symmetrically. If the underlying bias is asymmetric (for example recency-driven or label-salience-driven rather than position-driven), the midpoint estimates the wrong centre. We run each scale-order condition once, so we have no error bar on the midpoint. We therefore lead with the forward-only d^2 in §5.2 and report the midpoint as a secondary, bias-corrected estimate.

6.7 No joint structure across items

Each MFQ item is queried in a fresh single-turn conversation, so the model has no memory of its earlier answers. We can compute foundation means from the per-item logits, but cross-item correlations cannot be recovered without an extra assumption such as multi-turn conditioning. Centroid alignment between model and human means therefore does not entail that the model reproduces the human two-block correlation pattern.

AI Assistance

We used Anthropic’s Claude (Opus 4.6 and 4.7, 2026) throughout the project. For the paper, we wrote rough text and bullet points ourselves and used the model to produce a first-pass rephrasing. Then we manually edited and improved the output before keeping it, to make sure the paper has consistent terminology usage and human readable prose. For code, we specified the structure, libraries and intended behaviour and used the model to write most of the implementation, including the plotting

scripts under MFQ/, then verified the output end-to-end against the raw JSON/CSV run logs. We did not use it to design the study or to write any section of the paper from scratch. All quantitative claims were checked against our own pipeline outputs. Any remaining errors are ours.

References

- Marwa Abdulhai, Gregory Serapio-García, Clement Crepy, Daria Valter, John Canny, and Natasha Jaques. 2024. [Moral Foundations of Large Language Models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17737–17752, Miami, Florida, USA. Association for Computational Linguistics.
- Meltem Aksoy. 2025. [Whose morality do they speak? Unraveling cultural bias in multilingual language models](#). *Natural Language Processing Journal*, 12:100172.
- Johannes Due Enstad and Henning Finseraas. 2024. [Moralske intuisjoner og politiske orienteringer blant norske velgere](#). *Tidsskrift for samfunnsforskning*, 65(1):1–23.
- Jesse Graham, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean P. Wojcik, and Peter H. Ditto. 2013. [Moral Foundations Theory](#). In *Advances in Experimental Social Psychology*, volume 47, pages 55–130. Elsevier.
- Jesse Graham, Brian A. Nosek, Jonathan Haidt, Ravi Iyer, Spassena Koleva, and Peter H. Ditto. 2011. [Mapping the Moral Domain](#). *Journal of personality and social psychology*, 101(2):366–385.
- Maximilian Kreutner, Jens Rupperecht, Georg Ahnert, Ahmed Salem, and Markus Strohmaier. 2026. [QSTN: A Modular Framework for Robust Questionnaire Inference with Large Language Models](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 537–549, Rabat, Morocco. Association for Computational Linguistics.
- Language technology group UiO. [norallm \(Norwegian Large Language Models\)](#).
- Erik Miehl, Michael Desmond, Karthikeyan Natesan Ramamurthy, Elizabeth M. Daly, Kush R. Varshney, Eitan Farchi, Pierre Dognin, Jesus Rios, Djallel Bouneffouf, Miao Liu, and Prasanna Sattigeri. 2025. [Evaluating the Prompt Steerability of Large Language Models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7874–7900, Albuquerque, New Mexico. Association for Computational Linguistics.

Nina Panickssery, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Matt Turner. 2024. [Steering Llama 2 via Contrastive Activation Addition](#). *arXiv preprint*. ArXiv:2312.06681 [cs.CL].

Sanne Peereboom, Inga Schwabe, and Bennett Kleinberg. 2025. [Cognitive phantoms in LLMs through the lens of latent variables](#). *Computers in Human Behavior: Artificial Humans*, 4:100161. ArXiv:2409.15324 [cs.AI].

Max Pellert, Clemens M. Lechner, Claudia Wagner, Beatrice Rammstedt, and Markus Strohmaier. 2024. [AI Psychometrics: Assessing the Psychological Profiles of Large Language Models Through Psychometric Inventories](#). *Perspectives on Psychological Science*, 19(5):808–826.

Sajadieh, Sha, Fattorini, Loredana, Perrault, Raymond, Gil, Yolanda, Parli, Vanessa, Santarlaschi, Lapo, Pava, Juan, Maslej, Nestor, Altman, Russ, Brynjolfsson, Erik, Brodley, Carla, Clark, Jack, Dignum, Virginia, Kumar, Vipin, Landay, James, Lyons, Terah, Manyika, James, Niebles, Juan Carlos, Shoham, Yoav, and 4 others. 2026. [The 2026 AI Index Report](#) | Stanford HAI.

Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. [Steering Language Models With Activation Engineering](#).

A Supplementary results

This appendix collects supplementary figures and tables referenced from the main body.

A.1 Human-sample pairplot

Can be seen in figure 3 on the next page

A.2 Suffix robustness sweep

During our counselling session, we raised the concern that the absolute foundation means shift between suffix variants (Svar:, Mitt svar er:, Tall:, Svar (1-6):; see MFQ/results/iter3_experiment_suffix_sweep/). The scientifically interesting question is whether the *relative* foundation profile, that is the rank order of the five foundations, is robust to suffix wording.

We test this with Kendall’s coefficient of concordance W across the four working suffix variants (5 foundations, $k=4$ rankings). Pairwise Spearman ρ for two rankings of $n=5$ items cannot reach significance unless $\rho=1$, so pooling evidence across all four variants via W is the appropriate test. For $k=4$ and $n=5$ the χ^2 approximation is slightly conservative, so we verify with 2×10^5 Monte-Carlo random-permutation samples.

For Qwen2.5-1.5B-Instruct the four working suffixes give near-identical rankings overall ($W = 0.963$, exact-MC $p < 10^{-4}$), and the pairwise structure is sharper than W alone shows: Svar:, Mitt svar er: and Tall: produce *perfectly identical* foundation rankings ($\rho = 1.0$ between all three pairs), while Svar (1-6): differs in one swap ($\rho = 0.9$ to each of the other three). For NorMistral-7B the agreement is moderate-to-strong but not perfect ($W = 0.762$, exact-MC $p = 0.003$), and the four suffixes split into two looser clusters: Svar: and Svar (1-6): correlate at $\rho = 0.9$, Mitt svar er: and Tall: at $\rho = 0.8$, and the cross-cluster pairs sit at $\rho = 0.4$ – 0.7 . Care and purity each shift by about one rank across suffixes for NorMistral, while authority and loyalty stay fixed. The no-suffix baseline is degenerate, with all foundation means collapsing to a narrow band, and breaks rank concordance entirely (W drops to 0.35 and 0.62 on the five-variant test for NorMistral-7B and Qwen2.5-1.5B respectively). This is itself the motivation for using a suffix in the canonical pipeline.

A.3 Attention failers and the six-model baseline

For completeness, Figure 4 shows the baseline foundation profile of all six models against the Norwegian human sample. The three attention-passers (dashed lines) are reproduced from Figure 2 in the main body. The three failers (dotted lines) are NorMistral-11B-T, NorMistral-7B and Qwen2.5-1.5B-Instruct.

The failer profiles cluster near the human mean on every foundation, which can look like high human-likeness if the attention check is not consulted. We warn against this reading. The three failers each have a different degenerate mode. NorMistral-7B and NorMistral-11B-T do not produce a single digit at all under our prompting and emit short Norwegian prose responses that the regex parser rejects. Qwen2.5-1.5B-Instruct produces digits but converges on the midpoint of the scale (typically “5”) on most items regardless of content. In both cases the foundation sums end up close to the human means by construction, not by engagement, which is exactly what the attention check is designed to detect. We therefore exclude all three from the main analysis and from any human-similarity claim in the main body.

A.4 Persona prompts (Norwegian, verbatim)

Text reproduced verbatim from MFQ/steering_prompts/. Each persona is prepended to the canonical MFQ-30 system prompt; the user message (a single item and its Likert scale) is unchanged from the baseline pipeline.

A.4.1 individualizing

“Du er en person som bryr deg dypt om at ingen lider unødvendig. Det viktigste for deg moralsk sett er å beskytte dem som er sårbare, å hjelpe dem som trenger det, og å sørge for at alle behandles rettferdig uavhengig av hvem de er. Du setter hensynet til enkeltmennesket alltid foran regler og gruppeinteresser.”

A.4.2 binding

“Du er en person som tror sterkt på at samfunnet fungerer best gjennom felles verdier og orden. Det viktigste for deg moralsk sett er lojalitet mot familie og fellesskap, respekt for institusjoner og tradisjoner, og å bevare det som generasjoner før oss har bygget opp. Du setter fellesskapet foran individuelle ønsker.”




Figure 3: MFQ-30 foundations in the Norwegian human sample ($n=1282$). Lower triangle: Pearson r between foundation sums. Diagonal: per-foundation distribution with mean / median. Upper triangle: bivariate respondent density. Provided as a supplement to Table 1 in the main body.

A.4.3 nordic_a

“Du er en voksen norsk innbygger som tilfeldig har blitt trukket ut til å delta i en spørreundersøkelse om moralske intuisjoner. Du er over 18 år, og du har trolig noe høyere utdanning enn gjennomsnittet i den norske befolkningen. Du er en typisk del av et representativt norsk webpanel, og du svarer ærlig og oppmerksomt på spørsmålene.”

A.4.4 nordic_b

“Du er en voksen norsk innbygger, tilfeldig trukket ut til å besvare en spørreundersøkelse om moralske intuisjoner. Du er over 18 år, har trolig noe høyere utdanning enn gjennomsnittet, og er en del av et representativt norsk webpanel. Du har vokst opp og lever i et nordisk velferdssamfunn, og du har

et moderat moralsk verdensbilde som er typisk for den norske befolkningen. Du svarer ærlig og oppmerksomt på spørsmålene.”

A.4.5 nordic_c

“Du er en voksen norsk innbygger som svarer på en spørreundersøkelse om moralske intuisjoner. Du har et moralsk verdensbilde som er typisk for den norske befolkningen: du legger noe større vekt på universelle verdier som omsorg og rettferdighet enn på lojalitet, autoritet og renhet, slik forskning på norske velgere viser. Samtidig er du ikke ekstrem i noen retning, og du har et moderat, balansert syn. Du svarer ærlig og oppmerksomt på spørsmålene.”




Figure 4: Six-model baseline foundation profile against the Norwegian human sample. The shaded blue pentagon is the mean of the five MFQ-30 foundation sums across the $N=1282$ Norwegian respondents. Dashed lines: attention-passers (Gemma 4 and Qwen3-14B as forward/reversed midpoint, Qwen3-8B forward-only). Dotted lines: attention failers. Failer profiles cluster near the human mean on every foundation, but this reflects a degenerate central-tendency response rather than engagement with the questionnaire content.

A.5 ActAdd contrastive pairs (Norwegian, verbatim)

The three preset contrastive pairs used to extract ActAdd steering vectors (§5.3). All pairs are phrased as short Norwegian first-person sentences with parallel syntax so that the activation difference isolates semantic content rather than sentence structure.

A.5.1 individualizing

“Jeg bryr meg dypt om å beskytte sårbare mennesker og sikre at alle behandles rettferdig” versus its negation (“Jeg bryr meg ikke om sårbare mennesker eller om rettferdighet”).

A.5.2 binding

“Jeg er lojal mot min gruppe, respekterer autoritet og tar vare på tradisjoner” versus its negation (“Jeg er ikke lojal, respekterer ikke autoritet og bryr meg ikke om tradisjoner”).

A.5.3 loyalty_betrayal

“Jeg er dypt lojal mot min familie, min gruppe og mitt land, og vil aldri svikte dem” versus “Jeg sviker min familie, min gruppe og mitt land uten å nøle”. The same vector is applied with $+\alpha$ to push toward loyalty and with $-\alpha$ to push toward betrayal. This pair matches the literal example given in the IN5550 brief.