

Cultural Alignment of Open-Weight LLMs on the Inglehart–Welzel Map

Shav Vimalendiran
Sammy Labs, London, UK
shav@sammylabs.com

Abstract

Large language models increasingly mediate information and decisions across culturally diverse societies, raising questions about whose values their responses reflect. We examine whether geographic diversity among developers translates into cultural diversity in model responses. We administer the ten Integrated Values Surveys items defining the Inglehart–Welzel cultural map to 17 frontier open-weight models from Chinese and Western developers in English and Chinese, yielding 17,000 trials. An audited projection compares their responses with 109 countries and territories. On the primary map, 30 of 33 eligible model–language means are nearest a Western regional centroid. Under English administration, all 17 models give more permissive mean answers on homosexuality and abortion than the pooled human-survey mean. Chinese-language administration moves 15 of 16 paired models toward greater self-expression and increases both origin groups’ mean distance from the Confucian centroid; all 16 eligible Chinese-language means are nearest the Protestant European centroid. Origin differences nevertheless persist: Chinese-origin models are closer to the Confucian centroid than Western-origin models under English administration. The models thus exhibit predominantly Western-region proximity across developer origins, alongside differences within a restricted part of the cultural map.

1 Introduction

Culture shapes judgement, trust in automation, privacy expectations and health behaviour. As text generated by large language models (LLMs) enters these settings, culturally patterned outputs may affect users whose expectations differ from those represented in training and post-training. This motivates an empirical question: *what survey positions do models report under specified elicitation conditions?* A separate interpretive question is whether

their answers express a simulated personal position or estimate other people’s responses. The Inglehart–Welzel (IW) map summarises human values surveys along two axes, traditional versus secular-rational and survival versus self-expression (Inglehart and Welzel, 2005). Answers to its ten defining items can be projected onto the same plane as 109 mapped countries and territories; this does not establish that human and model answers measure the same underlying construct.

Our July 2024 public analysis (Vimalendiran, 2024) carried four defects in measurement code and elicitation design; our audit found three, and a reader of the released code reported the fourth. **(1) Fitting and projection defects:** the earlier code projected unstandardised model answers onto axes fitted to standardised survey data and re-fitted the rotation on model scores. Countries and models therefore did not share one frozen projection. The inherited missing-data loop also failed the claimed PPCA likelihood checks; we replace it with independently validated observed-data likelihood fitting (§3.2). **(2) An incorrectly prepared index:** the SPSS user-missing code `–3` on the autonomy index Y003 was treated as data in 31.9% of the original post-2005 source rows. EVS also omits the derived index while retaining its constituent answers; we now reconstruct valid indices before filtering and imputation. Its projection coefficient, approximately zero under the defective coding, is now (0.20, 0.31) (Appendix D). **(3) Pooled languages:** the one bilingual model’s runs were pooled despite a 1.87-map-unit English–Chinese gap. Map units denote Euclidean distance in the projected plane. The original analysis reported no uncertainty; most models had only one administration language, confounding language with developer origin. **(4) Rescaling offsets:** offsets carried over from Tao et al. (2024) were +0.38 and `–0.01`, rather than the World Values Survey (WVS) syntax’s +0.038 and `–0.10`. The correction translates the map by

(−0.342, −0.090), preserving pairwise distances but changing absolute coordinates (§3.2).

Contributions. (1) **Measurement.** An internally checked, unit-tested, open pipeline reconstructs the instrument from individual Integrated Values Surveys (IVS) data, combining the European Values Study (EVS) and WVS. Probabilistic principal component analysis retains the 30% of rows that listwise deletion would discard. Countries and models share one stored rotation; bootstrap summaries and alternative regional assignments expose selected sources of uncertainty. (2) **Reanalysis.** Eleven model–language combinations (cells) from 2024 are re-projected separately, revealing the bilingual model’s language contrast and correcting the original regional assignments. (3) **Experiment.** A crossed language-by-model design covers 17 models (10 Chinese-origin, 7 Western), 34 cells and 17,000 elicitation trials. Within-model contrasts separate the language comparison from developer-origin comparisons, while item-level and terminal-failure analyses qualify the geometric results (§4.2). (4) **Interpretation.** Five-annotator coding of 900 sampled reasoning excerpts identifies typicality or moderation targeting in a majority, with less agreement on persona reasoning. Prompt controls test the robustness of the projected positions, not the causal mechanism producing the explanations (§5). Throughout, a coordinate denotes a *reported survey position*, not a value demonstrated to be held by a model.

2 Related Work

Values-survey instruments. Arora et al. (2023) probe pre-trained models with WVS and Hofstede items. Tao et al. (2024) administer WVS items to GPT-family models, finding responses resembling English-speaking and Protestant European countries and showing that cultural prompting can shift them. Eren et al. (2026) extend this evaluation to open-weight models. Our emphasis is measurement auditability: individual-level IVS reconstruction, one projection shared by countries and models, and explicit uncertainty and sensitivity analyses under bilingual elicitation. Concurrently, Nguyen and Ahmad (2026) report prompt-tone shifts of up to 2.4 units on their IW reconstruction. LLM-GLOBE (Karinshak et al., 2024) finds China–US model differences on seven of nine dimensions in its open-generation analysis; both groups differ from their corresponding human pop-

ulations on most dimensions, which does not itself imply convergence between models. Luther and Brown (2025) find DeepSeek-V3 and V3.1 closer to the US than the Chinese VSM13 reference under all six prompting conditions they test, and Haslett et al. (2025) find ten Chinese and ten American models answering the WVS and MFQ-2.0 more like American than Chinese respondents, even under Chinese prompts or personas.

Elicitation-language effects. Bulté and Rigouts Terryn (2026) probe ten LLMs with Hofstede/WVS items in eleven languages, finding language-associated shifts alongside persistent proximity to a small set of countries. Effect sizes depend on the study: Sun and Wang (2025) report small overall language effects in a reassessment of Lu et al. (2025) using adapted tasks, newer models and more test items. These studies differ in models, instruments and design. We compare English–Chinese displacements within each model and ask whether those displacements differ by developer origin. Kazemi et al. (2024) associate online language resources with WVS representation error (44% of variance in GPT-4o; roughly fivefold greater errors for the lowest-resource languages). This motivates a training-data hypothesis but does not identify a causal mechanism; Mandarin is outside their main analysis. Maraia et al. (2026) find that sycophancy varies with register and authority, while models’ sycophancy profiles cluster by family across languages and regional variants. This is evidence about a different behaviour, not a direct explanation of our value-survey results.

Survey validity for LLM respondents. Rupprecht et al. (2026) run 334,800 simulated WVS interviews and document response biases, including model-dependent central tendencies and item non-response. Himelstein et al. (2026) show that refusal can mask value-laden bias in their experimental setting; this motivates inspecting missing responses here but does not identify our refusal mechanism. Taday Morochio et al. (2026) find no consistent aggregate benefit from persona conditioning in two models, with heterogeneous item and subgroup effects (see also Li et al., 2025). Our main protocol also uses a human persona, but specifies no nationality or demographic profile; its effects therefore require sensitivity checks rather than an assumption of neutrality.

Scope. We measure reported positions under a standardised survey-elicitation protocol, not deployment defaults or the full range of positions a model can express. The broader elicitation toolkit (Dang et al., 2026; Dang and Masud, 2026; Greco et al., 2026; Dai et al., 2025; Tan et al., 2026) and linguistic-relativity background (Au, 1983) are reviewed in Appendix I.

3 Method

3.1 Instrument and survey data

The IW map (Inglehart and Welzel, 2005) places societies on two axes (*survival vs. self-expression* and *traditional vs. secular-rational*) derived from ten IVS items, listed with their scales in Appendix B; the released code supplies the full bilingual prompts. We merge the EVS Trend File 1981–2017 (EVS, 2022) and the WVS Trend 1981–2022 (Haerpfer et al., 2022) with the published merge syntax. We retain respondents surveyed from 2005 onward (S020), recode out-of-range missing codes, and reconstruct absent Y003 from the four recorded child-quality indicators when all are valid 0/1 responses (Appendix D). Valid delivered indices are preserved. We then require at least six observed items per respondent. The fit uses 392,382 respondents from 112 country/territory codes. Country comparisons use the 109 mapped entities, covering 389,341 respondents; Northern Cyprus, Northern Ireland and Kosovo lack entries in the country-code table and are excluded after fitting. Country aggregation applies the IVS original national weight (S017). Survey years contribute according to their retained S017 weight mass within each country. Some residual missingness arises from items not fielded in particular country-waves, notably F118 (Appendix D). Model-based imputation retains the 30% of rows that complete-case analysis would discard, but its validity still depends on latent-model and missingness assumptions (Little and Rubin, 2019).

3.2 The projection pipeline

Probabilistic principal component analysis (PPCA; Tipping and Bishop, 1999) fits an unweighted two-dimensional Gaussian latent model, holding observed-item means and SDs fixed. We optimise its observed-data likelihood with L-BFGS-B, then complete missing entries with conditional means (Appendix D). The inherited approximate missing-data loop was replaced during the final audit. The retained scores explain 41.3% of completed-matrix

variance; conditional completion reduces that denominator.

Standardisation. Each item is z-scored with parameters *frozen at fit time*. The items natively span 1–2 (trust) to 1–10 (religiosity), so projecting *unstandardised* model answers onto axes fitted to *standardised* survey data (the 2024 defect) gave native item ranges and means inappropriate influence on the projected positions.

Rotation, fitted once and stored. A varimax rotation is fitted once, on the training score matrix, stored as an explicit orthogonal matrix, and thereafter only applied; $\hat{\theta} = -37.9^\circ$ is the solver’s output at tolerance 10^{-5} on a nearly flat criterion whose optimum lies 1.5° away. Fitting varimax to the score matrix with Kaiser row-normalisation is score-side, as in Rohe and Zeng (2023), but departs from their procedure (Appendix C). We adopt it because its output matches the recognised IW orientation, and report the full sensitivity grid: two of its six criteria place three or four cell means outside the quadrant. Because the rotation is applied to unwhitened scores, the two plotted axes are correlated at $r = 0.37$; no claim depends on their being uncorrelated.

Rescaling, and the frozen path. The WVS constants are $a = (1.81, 1.61)$ and $b = (0.038, -0.10)$, from the published syntax (Haerpfer et al., 2022). Correcting the earlier offsets translates the map by $(-0.342, -0.090)$; relative distances are preserved when references move with the points. These constants assume unit-variance factor scores. We divide rotated scores by their fitted standard deviations (SDs), $s = (1.47, 1.34)$, before applying the constants. Omitting this step inflates per-axis distances by those factors. The frozen path for any projected point is

$$\text{position} = \frac{((\mathbf{x} - \boldsymbol{\mu}) / \boldsymbol{\sigma}) CR}{s} a + b,$$

with every parameter fixed at fit time, read elementwise where shapes differ ($\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\sigma}$ are 10-vectors, C the 10×2 orthonormal projection axes, R the stored rotation, s, a, b per-axis pairs). We write the rescaled axes as $PC1'$ (survival–self-expression) and $PC2'$ (traditional–secular-rational). Distances throughout are Euclidean on this plane. This custom IVS-derived instrument is not externally calibrated to published IW country scores. One `project()` method enforces the shared transformation for countries and models. Three bundled

checks gate it: exact path identity over the 275,955 complete-case rows, correction accounting against our 2024 published coordinates ($R^2 = 0.988, 0.976$), and inspection of rotated projection coefficients, detailed in Appendix D.

3.3 Administering the survey to models

A *cell* is one model under one administration language. Each item is presented under ten persona prefixes, with five repeats: 500 elicitation trials per cell. The prefixes appear in the user message, not the system message. Six include an averaging cue (*average* or *typical*), three say only *a human being*, *a person* or *an individual*, and one says *a world citizen*. Appendix B lists the prefixes; the released `app/cloud_survey.py` supplies all item messages, format instructions and the system primer. A trailing system-role message then supplies the first-person prefill “Sure thing! Here is my numerical answer:”, kept from 2024; the Chinese counterpart is in Appendix B. Item-specific parsers validate responses; 2026 collection allows three attempts per pass plus transport re-sweeps (Limitations), versus 15 re-asks in 2024, and estimates are conditional on this protocol.

A cell enters the primary position analysis only if every item has at least ten parsed answers. The 2024 corpus contains ten models in eleven cells: seven English-only, two Chinese-only and one bilingual model. The 2026 design crosses 17 models with English and Chinese administration (the two *arms*, en and zh), yielding 34 cells and 17,000 recorded trials. Chinese wording was revised for 2026 (Appendix B). Within 2026 the language arms share the collection design except for F120’s anchor order (Limitations), but translation and inferred national context may contribute to their contrast. The 2024 records retain parsed values only. The 2026 records add final answer text, stored reasoning excerpts, variant and repeat IDs, attempt counts, the final attempt’s error and latency. Earlier failed-attempt texts are not retained. Respondent generation and thinking options were left at hosted defaults; their effective values and returned model revisions were not retained. All English runs preceded all Chinese runs on 4 August 2026, so administration time also varies with language.

3.4 Uncertainty

The reported point is the projection of the retained per-item means. We report bootstrap SDs and nominal confidence intervals (CIs). Because the

frozen projection is affine, arbitrary pairing of complete item responses into pseudo-respondents does not change this point. It can, however, change uncertainty estimates.

The *item bootstrap* ($B = 1,000$) resamples each item’s responses independently and omits cross-item covariance. It is not a guaranteed lower bound on uncertainty. The primary *cluster bootstrap* ($B = 10,000$) resamples the ten persona variants and pools their retained answers, preserving variant-level dependence. Empty variant–item groups contribute zero counts; only draws with no retained answer for an item use its observed mean as a fallback. The 2024 records lack variant IDs, so only the item bootstrap is available. Language contrasts use independent arm resampling; a matched-variant sensitivity is reported in Appendix E. Plotted ellipses use the replicate covariance and the normal-theory $\chi^2_{0.95,2}$ threshold. They are nominal regions for mean positions, not response distributions or calibrated 95% regions at ten clusters (Cameron et al., 2008).

Two descriptive region rules are reported: a radial-basis-function support-vector machine (RBF-SVM) and the nearest region centroid. The SVM’s selected-grid cross-validation score is 0.61; because tuning reuses the scoring folds, it is not an unbiased accuracy estimate. No headline depends on this classifier.

Classifier-free comparisons use the *survey reference* at (0.038, −0.10): the projection of the unweighted observed-item marginal-mean vector, equal to b by construction. The unweighted completed-respondent mean is (0.016, −0.104). We report reference-point distance, the fraction of 109 countries closer to the reference, and distance to the nearest non-Western region centroid. Here, non-Western excludes Protestant Europe, English-Speaking and Catholic Europe; the secular/self-expression quadrant lies above the reference on both axes. Alternative respondent-mean references are assessed separately. *Every replicate* describes the generated draws, not every possible resample or a simultaneous confidence guarantee (Appendix E).

3.5 Refusals, and the selection they induce

Refusals can induce informative missingness. In 2026, 215 of 17,000 trials failed to parse (1.26%; 83 en, 132 zh). The retained trial records contain 18,522 recorded attempts; exhausted transport calls need not enter that sum. Terminal failures include refusals and other invalid responses; they

concentrate on abortion (107) and homosexuality (82). The largest counts occur in Western models: gemma4:31b (28 English), nemotron-3-ultra (40 English, 88 Chinese) and gpt-oss:120b (10 English, 15 Chinese). The largest Chinese-origin count is minimax-m3: none in English and 19 in Chinese, including 13 on homosexuality (F118); qwen3.5:397b has three English and one Chinese failure. Refusal phrasing includes AI-identity boilerplate and premise rejection (Appendix B); marker-classified counts are conservative floors. Inclusion requires enough parsed responses on every item, and the most-refused items are the two largest positive self-expression projection coefficients. The sample is thus selected on a variable correlated with the outcome. Observed positions condition on adequate responses. For the 2024 nonresponding entries, filling missing answers with admissible extremes (a Manski-style bound; Manski, 2003) would not preserve an all-model quadrant claim (Limitations).

4 Results

Throughout, positions summarise the answers models report under this protocol. They are not direct measurements of internal values; the reasoning-text analysis in §5 examines possible response strategies.

4.1 The corrected 2024 cohort

The headline, classifier-free. Every 2024 cell mean (Figure 2) lies farther from the survey reference than at least 103/109 countries (94.5%, rounded). The minimum non-Western centroid distance is 1.61005849 map units. Across generated draws, the minima are 88/109 countries and 1.11702539 units: conservative floors are 80% and 1.11 units. Every draw remains in the quadrant. Six of the eleven cell means exceed Japan, the most secular country on the map ($PC2' = 1.83$). Confucian assignments change under the corrections, and the two region rules disagree on three cells (Table 1).

A large language contrast in one model. qwen2:7b is the only 2024 model administered in both languages. Its cells lie 1.87 map units apart, more than ten times its item-bootstrap axis SDs. Chinese administration yields more traditional responses on national pride, religiosity and respect for authority. The two originally Confucian-labelled cells lack matched English runs, so their developer origin and administration language are

confounded. The corrected SVM additionally labels qwen2:7b's English cell Confucian, whereas the nearest-centroid rule labels none Confucian.

4.2 The 2026 frontier generation, both arms

Observed response coverage. All ten Chinese-origin 2026 models yield usable cells in both languages. The recorded 2024 collection yielded usable output for four of nine Chinese-origin entries, or four of seven under a conservative adjustment for possible aliases (Appendix A). These convenience samples differ in model selection, serving, wording and retry protocol; their difference is descriptive, not a controlled estimate of progress over time. In 2026, every Chinese-origin cell parses at least 96.2% of trials. Terminal failures are concentrated in Western models, with minimax-m3 a Chinese-origin exception.

Positions. All 33 eligible cell means occupy the secular/self-expression quadrant (Figure 1); only nemotron-3-ultra [zh] is excluded. All 330,000 generated cluster-bootstrap replicates remain in that quadrant. The minimum cell-mean margin inside either boundary is 3.12 cluster SDs (Appendix E). The replicate coordinate minima are $PC1' = 0.27$ and $PC2' = -0.08$, against the survey reference (0.038, -0.10). Quadrant membership alone has limited specificity: it also holds for the all-midpoint vector ((0.36, 1.90)), 95.2% of uniform-random cell means and 29 of 109 country means. In contrast, only 29.5% of complete-case respondents occupy it, and none of 20,000 simulated 50-respondent cells reaches even the closest model's reference distance (Appendix I). Fourteen of seventeen English cells also fall inside the uniform baseline's central 95% distance band (1.62–2.45 units); neither quadrant membership nor reference distance alone uniquely identifies a cultural response strategy. Comparing each draw with the exact regenerated 2024 minima, 27 of 33 cells fall below either point threshold in some draw; 15 fall below either draw threshold. This post-review comparison uses full-precision artefacts (Appendix F). The closest cell, minimax-m3 [en], lies 1.63 units from the reference. Across every generated replicate of every eligible cell, at least 47 of 109 countries are closer to the reference and the nearest non-Western centroid is at least 0.37 units away (Appendix F).

Attenuation is arm-specific. English-arm reference distances range from 1.63 to 3.19 (median 2.24), compared with a median of 2.71 across the




Figure 1: The 2026 cohort, both arms: diamonds: English administration; triangles: Chinese administration; arrows en→zh per model; colour by origin; cluster-bootstrap nominal 95% regions. Fifteen of sixteen contrasts increase self-expression; secular movement is mixed. Dashed lines mark the survey reference; the cross marks the all-midpoint response vector. Labels show selected models and countries; Table 2 gives all cells.

eight 2024 English cells. Chinese-arm distances range from 2.15 to 3.31 (median 2.71). English-arm distances have a narrower interquartile range (IQR: 0.11 versus 0.36 map units), while mean pairwise distance is similar (0.82 versus 0.76). Thus, radial distances concentrate without comparable convergence on a single point (Appendix F). These are descriptive comparisons of different cohorts, scales, serving conditions and retry protocols (Limitations). Germany and Japan are the most frequent nearest countries in this two-dimensional projection. Some cells exceed Japan on secularity; none exceeds Sweden on self-expression (Appendix F).

Language of administration, now a designed factor. The within-model contrast $\delta_m = \bar{p}_{m,zh} - \bar{p}_{m,en}$ is estimable for 16 models, where $\bar{p}$ is the projected observed-item-mean position (Figure 6). The mean displacement magnitude is 0.64 map units (range 0.16–1.49; across-model bootstrap interval [0.51, 0.79]). The norm of

the mean displacement vector is 0.46. These intervals resample the 16 model estimates, hold within-model estimates fixed and treat releases as exchangeable; they are descriptive for this selected cohort. Family-weighted and matched-variant sensitivities appear in Appendix E. Signed components carry directional inference. The hypothesis that Chinese administration shifts positions towards the traditional pole is not supported: only 5 of 16 models move down the secular axis (two-sided sign $p = 0.21$; pre-specified directional $p = 0.96$). The mean secular shift is +0.18 (release-level interval [0.03, 0.33]), but the family-weighted interval spans zero (+0.10, [−0.10, 0.28]). 15 of 16 move towards self-expression, with mean shift +0.42 [0.22, 0.64]. Mean intervals and directional sign counts are different summaries of these heterogeneous contrasts. No origin-by-language interaction is detected: permutation tests give $p = 0.62$, 0.093 and 0.85 for $\delta PC1'$, $\delta PC2'$ and $\|\delta_m\|$. Mean

magnitudes are 0.63 for Chinese-origin and 0.66 for Western models. With only ten Chinese-origin and six Western paired models, including related releases, non-detection does not establish equivalence. The largest displacement belongs to *mistral-large-3:675b* (1.49 units [0.92, 2.00]).

Chinese administration does not strengthen Confucian proximity. Under English administration, Chinese-origin models lie nearer the Confucian centroid than Western models: 1.51 [1.39, 1.65] versus 1.87 [1.77, 1.98] (cluster-bootstrap CIs). Chinese administration increases these distances to 2.01 [1.90, 2.12] and 2.14 [2.04, 2.26], respectively. The Chinese-origin groups each contain ten models; the Western groups contain seven in English and six in Chinese after the exclusion. Three cells receive a Confucian label under both rules, all English-administered: *deepseek-v4-pro*, *minimax-m3* and *mistral-large-3:675b*. Under Chinese administration, nearest-centroid assigns every cell to Protestant Europe. These assignments describe sparse territory between the Protestant-European country cluster and Japan, not a full Confucian value profile. Across both languages, 30 of 33 eligible cell means are nearest a Western regional centroid. The language-induced-proximity hypothesis is not supported; the smaller Chinese-origin distance under English administration is a separate observation.

Where the language effect lives, item by item. Under English administration, all 17 models give more permissive mean answers on homosexuality (F118) and abortion (F120) than the pooled human-survey mean. Three item shifts survive Benjamini–Hochberg (BH) adjustment over ten items in exploratory release-level tests. We use two-sided exact binomial sign tests over the 17 models, dropping ties. Under Chinese administration, models report more interpersonal trust (A165: 15 of 16 non-tied models, BH $p = 0.003$), higher autonomy scores (Y003: 13 of 15, BH $p = 0.025$) and less petition-signing (E025: 16 of 17, BH $p = 0.003$). Omitting the four originally flagged Chinese autonomy comparisons removes its adjusted significance; traces suggest a study-skills reading in 11 of 15 trace-emitting models (Limitations). Averaging within nine developer families leaves no item below BH 0.05 (Appendix F). More permissive abortion answers (F120: 13 of 16, BH $p = 0.053$) and less respect for authority (E018: 13 of 17, BH $p = 0.098$) do not reach the adjusted

threshold; neither do the other five items. The tests include *nemotron-3-ultra*’s Chinese cell, whose abortion mean uses four answers. Excluding this model changes F120’s adjusted p to 0.088 but leaves the same three items BH-significant (Appendix F). The predicted traditional-ward shifts on national pride (G006), religiosity (F063) and authority (E018) are not supported individually. Petition-signing explicitly asks about participation in a political action. Political-context sensitivity could explain its contrast with trust and autonomy, but language, translation and inferred national context change together. Non-detection on other items does not establish unchanged values. In a post-hoc contrast, petition-signing’s shift is more survival-ward than the mean of the other four refusal-sensitive items in 16 of 17 models (exact sign $p = 0.0003$, descriptive; Appendix F).

Refusal is language-dependent, with model-specific sign. Refusal changes have heterogeneous directions: two opposite-sign model swings dominate, and a paired sign test detects no aggregate direction ($p = 0.29$; Appendix F). Separately minimising each projected coordinate over admissible fills of terminal failures leaves both minima inside the quadrant for every 2026 cell. Even *nemotron-3-ultra* [zh] remains inside ($PC1' = 0.37$, $PC2' = 0.77$, each from its own adversarial fill; Appendix F). This bound concerns terminal failures after retries, not selection introduced by the retry process itself. Replacing all F118/F120 answers with human item means moves three cells to or past the reference on $PC1'$ (Appendix C). Quadrant placement, not only distance, is therefore sensitive to these items.

Descriptive sources of dispersion. In a sums-of-squares partition of constructed profiles (Table 4), models account for 21.4% and 22.5% of self-expression and secular variation, and language within model for 10.7% and 6.4%; sibling contrasts are descriptive (Appendix C).

5 Discussion

Reported positions and verbal explanations. Five LLM annotators coded 900 stratified reasoning excerpts from 30 trace-emitting cells; four cells from two non-emitting models are absent. Majority votes identify typicality or moderation targeting in 55% (rater range 50–65%), persona reasoning in 5% (2–17%) and AI-identity or guideline ref-

erences in 42% (35–48%). Fleiss’ κ is respectively 0.78, 0.35 and 0.75; the persona estimate has only fair agreement. The first category includes middle-of-the-road answers, not just estimates of population typicality. Moreover, 164 excerpts reach the 2,000-character cap, and two annotators each rate 60 own-model excerpts (Appendix H). These overlapping labels describe sampled text, not necessarily faithful causal explanations (Chen et al., 2025b). An independent review also found positive labels on role restatements followed by arbitrary picks (Appendix H). The broad category and these coding errors limit what can be inferred about population-statistic estimation or held values.

Several explanations remain compatible with the geometry. Training-data composition, post-training preferences, response-scale effects and selective answering could all contribute. The three 2024 dolphin fine-tunes remain near the main cluster, but without matched interventions this does not isolate any training stage. Scale geometry matters: a vector of item-scale midpoints projects to (0.36, 1.90), above Japan on the secular axis. In contrast, the observed item-marginal modes project to (−1.47, −1.47), under either unweighted or S017-weighted counting. The vector of marginal modes need not equal the joint mode. Typicality-coded explanations therefore do not demonstrate accurate recovery of human survey modes. Raw-item midpoint distance has little rank association with secularity (Spearman $\rho = 0.02$), without excluding a common offset. Neither that association nor reference-point distance identifies accurate typicality: population means, modes and scale midpoints are different targets. The language-dependent failures we observe leave selection as a possible contributor; psychometric nonresponse is also documented elsewhere (Xie et al., 2025).

Release differences do not identify training causes. Prior work shows that post-training can alter reported opinions (Santurkar et al., 2023; Ryan et al., 2024; Rozado, 2024; Choenni et al., 2024); preference construction itself involves human choices (Ouyang et al., 2022; Huang et al., 2024). Our between-release comparisons do not isolate those choices from pretraining, architecture or serving. On four items (F118, F120, Y003, E025), all 17 English-elicited models deviate from the human mean in the same direction. Little net preference for higher response codes rules out neither item-specific acquiescence nor answer-order effects.

Shared training histories and synthetic data could make model releases statistically dependent (Sun et al., 2025; Shimabucoro et al., 2024; Gudibande et al., 2023; Jiang et al., 2025); the present design does not identify that pathway. Survey-standardised profiles are more similar within than across origin groups, an association sensitive to the comparison unit and not a causal origin effect (Appendix I).

Prompt controls support the geometry, not a reasoning mechanism. Removing an averaging cue or the persona prefix does not eliminate the quadrant pattern among eligible cells (Appendix G). With the three non-averaging prefixes and at least one answer per item, all 32 estimable cells remain in the quadrant; 29 are farther from the reference. Requiring at least ten answers per item, all 13 eligible no-persona cells (of 17) remain in the quadrant and 12 are farther away, chiefly through greater secularity. Median distance is 3.43 versus 2.21 in their matched primary cells. The latter condition retains the first-person system primer and was collected five weeks later; hosted-snapshot or serving changes may contribute. Existing majority-vote labels do show a prefix association: typicality-or-moderation codes occur in 79.6% of selected English excerpts with averaging cues versus 33.8% with bare prefixes (Chinese: 56.7% versus 31.1%; Appendix H). These selected, non-independent excerpts do not identify a causal effect, and link framing to the reasoning-text pattern.

Scope of the result. The models concentrate in one part of the map: 31 of 33 means occupy the top quartile of both country-coordinate distributions. Within that area, model and prompt differences remain substantial. The quadrant result describes pooled survey responses under the stated instrument and protocol; it does not establish full cultural profiles, deployment defaults or invariant traits. Individual prompt variants can lie outside the quadrant despite precise pooled means. Prior work likewise finds condition- and model-dependent stability (Moore et al., 2024; Kovač et al., 2024; Huang, 2026), while prompting can alter positions (Tao et al., 2024; Dang et al., 2026; Dang and Masud, 2026). Deployment therefore requires testing elicitation sensitivity and downstream behaviour in the intended setting.

Limitations

Conditional on retained answers. The primary analysis requires at least ten parseable answers to each item. The most-refused items are the two largest positive self-expression projection coefficients, so selection can affect the measured outcome (§3.5). Worst-case imputation of terminal failures preserves the 2026 quadrant result, including the excluded cell, but not all distance claims (§4.2). This bound applies to the implemented retry protocol, not to unrecorded rejected drafts or hypothetical first-response distributions. The sparse 2024 failures do not support the same unconditional inference.

Retry and prompt implementation. The 2024 protocol allowed up to 15 re-asks. In 2026 each collection pass allows three attempts per trial, with at most five passes per collector run for trials deferred by transport failures: up to 15 calls per trial per run, and no bound across resumed runs. During collection this included trials with earlier parse failures, so refusing trials on unreliable endpoints could receive more chances to answer. Retaining the first parseable response conditions the distribution on numerical compliance. The 2024 corpus lacks attempt counts; 2026 records attempts per pass, not every rejected draft or re-sweep. The primer is a trailing system message holding a first-person answer prefill, whose rendering depends on the chat template; the no-persona control retains it. Persona prefixes are in the user message. Chinese wording and instructions also changed between cohorts; identical instrument items do not imply identical prompts or serving conditions.

Elicitation is not neutral. Models may respond to an inferred interlocutor; political-bias audits can reflect sycophancy towards the perceived auditor (Törnberg and Schimmel, 2026). Our human-persona framing is an explicit elicitation choice. The non-averaging and no-persona checks preserve the geometric pattern among eligible cells (§5, Appendix G), but do not test its reasoning mechanism. The latter is a later, hosted collection retaining the system primer, not a no-system-prompt experiment. Named-auditor and nationality-crossed controls remain untested.

Construct and trace validity. Model survey answers need not predict open-ended behaviour; human survey-behaviour correlations do not transfer

automatically. Midpoint responding partly mimics secularity on this instrument. The trace audit covers selected, sometimes truncated explanations and uses LLM annotators, including two self-raters. Agreement on the broad typicality-or-moderation code does not establish actual population-statistic estimation or reasoning faithfulness; persona coding has only fair agreement. Coordinates are therefore reported survey positions, not demonstrated internal values.

Translation and inferred referents. The Chinese prompts are our translations, not the official WVS Chinese questionnaire. The Chinese Y003 stem is ambiguous between qualities children learn at home and studying at home: a prompt-stripped lexical scan flags study-related wording in 11 of the 15 trace-emitting models, three of them Western, with examples supporting that reading. The four-model exclusion in Appendix F is illustrative; the language contrast cannot isolate autonomy from this wording ambiguity (§4.2). The English and Chinese F120 prompts also name the scale anchors in opposite order, so F120's language contrast is confounded with anchor order. G006 and Y002 leave “your nationality” and “this country” unspecified. Language may change the inferred referent as well as the interpreted proposition. Appendix B documents earlier translation corrections. Exclusion sensitivity does not identify the effect of a corrected translation.

Instrument familiarity and dated reference. The public IVS items may occur in training data. At least one sampled explanation names the WVS and discusses country response ranges (Appendix I); remembered survey information may influence answers. The projection is frozen on respondents from 2005 onward in the EVS/WVS trend releases, not a survey of human values in 2026. Its reference point projects observed-item marginal means, rather than representing the mean completed respondent or a population-weighted global distribution.

Unpropagated instrument uncertainty. Country coordinates inherit imputation uncertainty that is not included in model sampling or prompt-cluster intervals. Country-point precision is therefore not directly comparable with the plotted model regions (Figures 1 and 2). We reconstruct valid autonomy indices from recorded constituents before imputing residual missing entries. Remaining country-item gaps require model assumptions, even when miss-

ingness is by design. Internal checks do not externally calibrate this unweighted, restricted-period PPCA and score-rotation map to published WVS country scores.

Small-cluster uncertainty. The conventional 95% χ^2 ellipses have approximately 85% coverage in a Gaussian ten-cluster simulation; actual prompt-distribution coverage is unknown (Appendix E). Ten-cluster bootstrap intervals and the absence of violations in finite generated replicates should not be read as calibrated simultaneous confidence guarantees. Model-level summaries additionally depend on how variation across related releases is treated.

Rotation and seed sensitivity. Two of the six rotation criteria in Table 3 place three or four cell means outside the quadrant, so the headline is conditional on the chosen criterion. Across 20 likelihood-fitting seeds, the rotation angle spans approximately 0.000016° and the maximum country-coordinate range is below 10^{-6} map units; this bounds likelihood-fit initialisation only, not sampling, missingness assumptions or the varimax solver, whose stored solution lies 1.5° from the criterion optimum (Appendix C).

Convenience cohorts, not a longitudinal panel. The cohorts share no models and differ in recorded serving conditions (local versus cloud), model scales, retries and inclusion. Effective cloud sampling settings and immutable returned model revisions were not retained; English collection preceded Chinese collection. Historical aliases further complicate counts of distinct attempted weights. Usable-corpus rates are descriptive comparisons of recorded convenience samples, not evidence of generational improvement. Neither cohort is a random sample of open-weight models, and related releases are not independent draws. Population-wide and causal claims about developer origin or training stages are unsupported.

Ethics Statement

This work measures reported model survey positions using publicly documented instruments. No new human participants were recruited; the pre-existing IVS collections are used under their data-use agreements. We do not redistribute the IVS microdata. Survey item texts are © WVS/EVS and are reproduced for research and replication. Cultural-region labels are the IW map’s analytical categories, not judgements of societies. Proximity on two axes

does not describe an entire culture, establish population representativeness or show which values are correct. Steering towards any cultural profile is a normative choice that this measurement study does not endorse. Model-generated explanations and LLM coding are not treated as evidence of conscious or internally held values.

Acknowledgements

We thank the two anonymous ORACLE reviewers, whose requests motivated the prompt controls and multi-annotator reliability analysis in Appendices G and H, and the workshop organisers. We thank the GitHub user kwinkunks, who identified the rescaling-offset error in the released code (§3.2). The World Values Survey Association, the European Values Study and GESIS provide the survey data; the pipeline retains interface code adapted from `pca-magic` (Apache-2.0), with its fitting loop replaced. The author used Claude (Anthropic) and Codex (OpenAI) for coding, analysis checks, reference verification and drafting, and reviewed the resulting work. The release documents inputs and scripts for the primary analyses; IVS microdata must be obtained separately under the applicable data-use agreements.

References

- Emmanuel Ameisen, Jack Lindsey, Adam Pearce, Wes Gurnee, Nicholas L. Turner, Brian Chen, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermy, Andy Jones, and 8 others. 2025. [Circuit tracing: Revealing computational graphs in language models](#). *Transformer Circuits Thread*.
- Arnav Arora, Lucie-Aimée Kaffee, and Isabelle Augenstein. 2023. [Probing pre-trained language models for cross-cultural differences in values](#). In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 114–130, Dubrovnik, Croatia. Association for Computational Linguistics.
- Terry Kit-Fong Au. 1983. [Chinese and English counterfactuals: The Sapir-Whorf hypothesis revisited](#). *Cognition*, 15(1–3):155–187.
- Bram Bulté and Ayla Rigouts Terryn. 2026. [LLMs and cultural values: The impact of prompt language and explicit cultural framing](#). *Computational Linguistics*, 52(2):407–494. arXiv:2511.03980.
- Maarten Buyt, Alexander Rogiers, Sander Noels, Guillaume Bied, Iris Dominguez-Catena, Edith Heiter,

- Iman Johary, Alexandru-Cristian Mara, Raphaël Romero, Jeffrey Lijffijt, and Tijl De Bie. 2026. [Large language models reflect the ideology of their creators](#). *npj Artificial Intelligence*, 2(1):7. arXiv:2410.18417.
- A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. 2008. [Bootstrap-based improvements for inference with clustered errors](#). *The Review of Economics and Statistics*, 90(3):414–427.
- David Chanin, James Wilken-Smith, Tomáš Dulka, Hardik Bhatnagar, Satvik Golechha, and Joseph Bloom. 2025. [A is for absorption: Studying feature splitting and absorption in sparse autoencoders](#). In *Advances in Neural Information Processing Systems (NeurIPS 2025)*. Oral. arXiv:2409.14507.
- Runjin Chen, Andy Ardit, Henry Sleight, Owain Evans, and Jack Lindsey. 2025a. [Persona vectors: Monitoring and controlling character traits in language models](#). arXiv preprint. arXiv:2507.21509.
- Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, Vlad Mikulik, Samuel R. Bowman, Jan Leike, Jared Kaplan, and Ethan Perez. 2025b. [Reasoning models don’t always say what they think](#). arXiv preprint. arXiv:2505.05410.
- Rochelle Choenni, Anne Lauscher, and Ekaterina Shutova. 2024. [The echoes of multilinguality: Tracing cultural value shifts during language model fine-tuning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*, pages 15042–15058. arXiv:2405.12744.
- Xunlian Dai, Li Zhou, Benyou Wang, and Haizhou Li. 2025. [From word to world: Evaluate and mitigate culture bias in LLMs via word association test](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP 2025)*, pages 24510–24526. Oral. arXiv:2505.18562.
- Trung Duc Anh Dang, Tung Kieu, and Sarah Masud. 2026. [Scenario-based probing and steering cultural values in large language models — extended version](#). arXiv preprint. arXiv:2606.11399.
- Trung Duc Anh Dang and Sarah Masud. 2026. [Cultural value alignment via latent activation steering in large language models](#). arXiv preprint. arXiv:2605.26365. Presented at the ACL 2026 Student Research Workshop (non-archival track).
- Joshua Engels, Logan Smith, and Max Tegmark. 2025. [Decomposing the dark matter of sparse autoencoders](#). *Transactions on Machine Learning Research*.
- Maksim E. Eren, Eric Michalak, Brian Cook, and Johnny Seales, Jr. 2026. [Prompt programming for cultural bias and alignment of large language models](#). In *Proceedings of the 2026 ACM Symposium on Document Engineering*, pages 1–10. ACM. arXiv:2603.16827.
- EVS. 2022. [EVS trend file 1981–2017](#). GESIS Data Archive, Cologne. ZA7503 Data file Version 3.0.0.
- Candida M. Greco, Lucio La Cava, and Andrea Tagarelli. 2026. [Culturally grounded personas in large language models: Characterization and alignment with socio-psychological value frameworks](#). arXiv preprint. arXiv:2601.22396.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. [The false promise of imitating proprietary LLMs](#). arXiv preprint. arXiv:2305.15717.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. [DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning](#). *Nature*, 645(8081):633–638. Distillation provenance is addressed in the supplementary materials.
- Christian Haerpfer, Ronald Inglehart, Alejandro Moreno, Christian Welzel, Kseniya Kizilova, Jaime Diez-Medrano, Marta Lagos, Pippa Norris, Eduard Ponarin, and Bi Puranen. 2022. [World values survey trend file \(1981–2022\) cross-national data-set](#). Madrid, Spain & Vienna, Austria: JD Systems Institute & WWSA Secretariat. Credited contributors are dataset editors. Data File Version 4.0.0 (2024-06-30); author-reported retrieval 11 August 2024.
- Jongwook Han, Jongwon Lim, Injin Kong, and Yohan Jo. 2026. [Dual mechanisms of value expression: Intrinsic vs. prompted values in large language models](#). In *Proceedings of the International Conference on Machine Learning (ICML 2026)*. arXiv:2509.24319.
- David Haslett, Linus Ta-Lun Huang, Leila Khalatbari, Janet Hui-wen Hsiao, and Antoni B. Chan. 2025. [Made-in China, thinking in America: U.S. values persist in Chinese LLMs](#). arXiv preprint. arXiv:2512.13723.
- Rom Himelstein, Amit LeVi, Brit Youngmann, Yaniv Nemcovsky, and Avi Mendelson. 2026. [Silenced biases: The dark side LLMs learned to refuse](#). In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2026)*, *AI Alignment track*. Oral. Preprint 2025, arXiv:2511.03369.
- Haonan Huang. 2026. [The yes-no bias of large language models reflects answer order and wording, not shifts in moral judgment](#). arXiv preprint. arXiv:2607.05552.
- Saffron Huang, Divya Siddarth, Liane Lovitt, Thomas I. Liao, Esin Durmus, Alex Tamkin, and Deep Ganguli. 2024. [Collective constitutional AI: Aligning a language model with public input](#). In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT 2024)*, pages 1395–1417.
- Ronald Inglehart and Christian Welzel. 2005. [Modernization, Cultural Change, and Democracy: The Human Development Sequence](#). Cambridge University Press.

- Liwei Jiang, Yuanjun Chai, Margaret Li, Mickel Liu, Raymond Fok, Nouha Dziri, Yulia Tsvetkov, Maarten Sap, and Yejin Choi. 2025. [Artificial hivemind: The open-ended homogeneity of language models \(and beyond\)](#). In *Advances in Neural Information Processing Systems (NeurIPS 2025), Datasets and Benchmarks Track*. NeurIPS 2025 proceedings version.
- Elise Karinshak, Amanda Hu, Kewen Kong, Vishwanatha Rao, Jingren Wang, Jindong Wang, and Yi Zeng. 2024. [LLM-GLOBE: A benchmark evaluating the cultural values embedded in LLM output](#). arXiv preprint. arXiv:2411.06032.
- Sharif Kazemi, Gloria Gerhardt, Jonty Katz, Caroline Ida Kuria, Estelle Pan, and Umang Prabhakar. 2024. [Cultural fidelity in large-language models: An evaluation of online language resources as a driver of model performance in value representation](#). arXiv preprint. arXiv:2410.10489.
- Simran Khanuja, Hongbin Liu, Shujian Zhang, John Lambert, Mingqing Chen, Rajiv Mathews, and Lun Wang. 2026. [Steering LLMs for culturally localized generation](#). arXiv preprint. arXiv:2603.23301.
- Grgur Kovač, Rémy Portelas, Masataka Sawayama, Peter Ford Dominey, and Pierre-Yves Oudeyer. 2024. [Stick to your role! Stability of personal values expressed in large language models](#). *PLOS ONE*, 19(8):e0309114. arXiv:2402.14846.
- Thom Lake, Eunsol Choi, and Greg Durrett. 2025. [From distributional to Overton pluralism: Investigating large language model alignment](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL 2025)*, pages 6794–6814. arXiv:2406.17692.
- Dai Li, Linzhuo Li, and Huilian Sophie Qiu. 2025. [ChatGPT is not A Man but Das Man: Representativeness and structural consistency of silicon samples generated by large language models](#). arXiv preprint. arXiv:2507.02919.
- Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermy, Andy Jones, and 8 others. 2025. [On the biology of a large language model](#). *Transformer Circuits Thread*.
- Roderick Little and Donald Rubin. 2019. *Statistical Analysis with Missing Data*, 3rd edition. Wiley.
- Jackson G. Lu, Lesley Luyang Song, and Lu Doris Zhang. 2025. [Cultural tendencies in generative AI](#). *Nature Human Behaviour*, 9(11):2360–2369.
- James Luther and Donald Brown. 2025. [DeepSeek’s WEIRD behavior: The cultural alignment of large language models and the effects of prompt language and cultural prompting](#). arXiv preprint. arXiv:2512.09772.
- Charles F. Manski. 2003. *Partial Identification of Probability Distributions*. Springer.
- Gabriele Maraia, Fabio Massimo Zanzotto, and Leonardo Ranaldi. 2026. [Sounding vs. being an expert: Disentangling authority, register and cultural impact in sycophantic LLMs](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 32492–32508, San Diego, California, United States. Association for Computational Linguistics.
- Jared Moore, Tanvi Deshpande, and Diyi Yang. 2024. [Are large language models consistent over value-laden questions?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15185–15221. arXiv:2407.02996.
- An Duy Nguyen and Muhammad Aurangzeb Ahmad. 2026. [Measurement validity in LLM cultural alignment](#). arXiv preprint. arXiv:2608.29266.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems (NeurIPS 2022)*. arXiv:2203.02155.
- Karl Rohe and Muzhe Zeng. 2023. [Vintage factor analysis with Varimax performs statistical inference](#). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 85(4):1037–1060.
- David Rozado. 2024. [The political preferences of LLMs](#). *PLOS ONE*, 19(7):e0306621.
- Jens Rupperecht, Georg Ahnert, and Markus Strohmaier. 2026. [Prompt perturbations reveal human-like biases in large language model survey responses](#). In *Proceedings of the Seventh Workshop on Natural Language Processing and Computational Social Science (NLP+CSS 2026)*, pages 1–21, San Diego, California, United States. Association for Computational Linguistics. arXiv:2507.07188.
- Michael J. Ryan, William Held, and Diyi Yang. 2024. [Unintended impacts of LLM alignment on global representation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*, pages 16121–16140. arXiv:2402.15018.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. [Whose opinions do language models reflect?](#) In *Proceedings of the International Conference on Machine Learning (ICML 2023)*, volume 202 of *Proceedings of Machine Learning Research*, pages 29971–30004. arXiv:2303.17548.
- Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph

- Bloom, Stella Biderman, Adria Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, William Saunders, and 10 others. 2025. [Open problems in mechanistic interpretability](#). *Transactions on Machine Learning Research*.
- Luísa Shimabucoro, Sebastian Ruder, Julia Kreutzer, Marzieh Fadaee, and Sara Hooker. 2024. [LLM see, LLM do: Leveraging active inheritance to target non-differentiable objectives](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*, pages 9243–9267. arXiv:2407.01490.
- Kun Sun and Rong Wang. 2025. [The fragility of “cultural tendencies” in LLMs](#). arXiv preprint. arXiv:2510.05869.
- Mingjie Sun, Yida Yin, Zhiqiu Xu, J Zico Kolter, and Zhuang Liu. 2025. [Idiosyncrasies in large language models](#). In *Proceedings of the International Conference on Machine Learning (ICML 2025)*, volume 267 of *Proceedings of Machine Learning Research*, pages 57854–57885. arXiv:2502.12150.
- Erika Elizabeth Taday Morocho, Lorenzo Cima, Tiziano Fagni, Marco Avvenuti, and Stefano Cresci. 2026. [Assessing the reliability of persona-conditioned LLMs as synthetic survey respondents](#). In *Companion Proceedings of the ACM Web Conference 2026*, pages 320–329. arXiv:2602.18462.
- Qian Tan, Lei Jiang, Yuting Zeng, Shuoyang Ding, and Xiaohua Xu. 2026. [Mitigating cultural bias in LLMs via multi-agent cultural debate](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 8600–8612, San Diego, California, United States. Association for Computational Linguistics. arXiv:2601.12091.
- Yan Tao, Olga Viberg, Ryan S Baker, and René F Kizilcec. 2024. [Cultural bias and cultural alignment of large language models](#). *PNAS Nexus*, 3(9):pgae346. Preprint: arXiv:2311.14096.
- Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, Alex Tamkin, Esin Durmus, Tristan Hume, Francesco Mosconi, C. Daniel Freeman, and 7 others. 2024. [Scaling monosemanticity: Extracting interpretable features from Claude 3 Sonnet](#). *Transformer Circuits Thread*.
- Michael E. Tipping and Christopher M. Bishop. 1999. [Probabilistic principal component analysis](#). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622.
- Petter Törnberg and Michelle Schimmel. 2026. [Political bias audits of LLMs capture sycophancy to the inferred auditor](#). arXiv preprint. arXiv:2604.27633.
- Allen Tran. 2015. [pca-magic](#). GitHub repository. Apache-2.0 licensed.
- Veniamin Veselovsky, Berke Argın, Benedikt Stroebl, Chris Wendler, Robert West, James Evans, Thomas L. Griffiths, and Arvind Narayanan. 2026. [Localized cultural knowledge is conserved and controllable in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 43152–43178. arXiv:2504.10191.
- Shav Vimalendiran. 2024. [Cultural bias in LLMs](#). Blog post, shav.dev, 20 July 2024. Superseded and corrected by this paper.
- Zhengxuan Wu, Aryaman Arora, Atticus Geiger, Zheng Wang, Jing Huang, Dan Jurafsky, Christopher D Manning, and Christopher Potts. 2025. [AxBench: Steering LLMs? Even simple baselines outperform sparse autoencoders](#). In *Proceedings of the International Conference on Machine Learning (ICML 2025)*, volume 267 of *Proceedings of Machine Learning Research*, pages 67035–67080. arXiv:2501.17148.
- Wei Xie, Shuoyoucheng Ma, Zhenhua Wang, Xiaobing Sun, Kai Chen, Enze Wang, Wei Liu, and Hanying Tong. 2025. [AIPsychoBench: Understanding the psychometric differences between LLMs and humans](#). In *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci 2025)*, pages 87–94.
- Taisei Yamamoto, Ryoma Kumon, Danushka Bollegala, and Hitomi Yanaka. 2026. [Neuron-level analysis of cultural understanding in large language models](#). In *Proceedings of the International Conference on Learning Representations (ICLR 2026)*. arXiv:2510.08284.

A Models

2024 cohort (eleven model-language cells). Positions are in Table 1. Author records describe local Ollama serving with Q4-quantised community GGUF builds on consumer hardware; immutable per-run weight snapshots and hardware logs were not retained.

2024 entries recorded as producing no usable corpus. yi:34b, aquilachat2:34b, glm4:9b, xuanyuan:70b, kingzeus/llama-3-chinese-8b-instruct-v3 (five recorded model entries; distinct-weight provenance is qualified below). Several were not then distributed through the Ollama registry and were deployed via hand-written Modelfiles wrapping community GGUF quantisations; that workaround is why deployment provenance for these runs is thinner than for the responsive cohort, and it is one reason the 2026 protocol moved to cloud serving with trial identifiers and terminal-response metadata. Hosted weight digests remain

unresolved for some tags. Contemporaneous author notes describe the following failure modes; the failed raw outputs were not retained and these descriptions cannot be independently rechecked: yi:34b and glm4:9b answered with a bare “.”; aquilachat2:34b answered 。 or echoed the prompt; xuanyuan:70b produced unintelligible output; kingzeus/...-v3 failed intermittently.

2026 cohort. Seventeen cloud-served models. Chinese-origin: deepseek-v4-flash, deepseek-v4-flash:0731, deepseek-v4-pro, glm-5.1, glm-5.2, kimi-k2.6, kimi-k2.7-code, minimax-m2.7, minimax-m3, qwen3.5:397b. Western: gemma4:31b, gpt-oss:20b, gpt-oss:120b, mistral-large-3:675b, nemotron-3-nano:30b, nemotron-3-super, nemotron-3-ultra. Effective serving precision was not retained in the collection records and remains an unresolved difference between the cohorts.

nemotron-3-ultra [zh] is excluded from position estimates (F120 parsed 4/50, under the 10-response threshold; its worst-case Manski bound still lies inside the secular/self-expression quadrant, §4.2), so its δ_m is not estimable; all other 33 cells meet the threshold on every item, and no inclusion decision changes at thresholds 5 or 25 (Appendix C). Its marker-classified refusal count is a floor: every one of that cell’s 34 residual format-classified failures is a first-person prose declination verified by reading each, so its substantive refusals lie between 54 and 88.

Usable-corpus rates: 4/9 (Clopper–Pearson 95% CI [14%,79%]) recorded Chinese-origin entries produced a usable corpus in 2024 against 10/10 in 2026 (95% CI [69%,100%]). These denominators count collected model entries; pre-collection availability exclusions are recorded in the released protocol. The 4/9 takes the recorded 2024 model list at face value; the yi and glm Modelfiles both point at the AquilaChat2 GGUF, so three of the five failures may have been one model tested three times, which would put the denominator as low as seven (4/7, Clopper–Pearson 95% CI [18%, 90%]). Fisher’s exact two-sided test on the contrast gives $p = 0.011$ against 4/9 and $p = 0.051$ against 4/7. These tests compare the recorded convenience cohorts; they do not identify a generational improvement. For this statistic “Chinese-origin” means Chinese-developed or Chinese fine-tuned in 2024, and Chinese-developed in 2026; the denomina-

tors are convenience samples of what was locally runnable or cloud-served, not a fixed population.

Region-rule disagreement in 2024. The SVM assigns three cells to the Confucian region while nearest-centroid assigns *all eleven* to Protestant Europe; the disagreement is informative (the three calls extrapolate into map territory containing no countries, their nearest country Japan, an outlier within its own region), and one of the three is the *English-administered* qwen2:7b, reported as Protestant European in the 2024 analysis: the Confucian region is not even stable across the corrections.

B The ten IVS items, the elicitation protocol, and refusal taxonomy

Items. The released app/cloud_survey.py supplies complete item messages, format instructions, prefix variants and the system primer in both languages. A008 happiness (1–4), A165 trust (1–2), E018 respect for authority (1–3), E025 petition (1–3), F063 importance of God (1–10), F118 justifiability of homosexuality (1–10), F120 justifiability of abortion (1–10), G006 national pride (1–4), Y002 post-materialist index (two ranked goals from four), Y003 autonomy index (up to five child qualities from eleven; official syntax $Y003 = (A029 + A039) - (A040 + A042)$, independence plus determination minus religious faith and obedience, on 0/1 mentioned codings).

Prompt variants. The ten English prefixes, each followed by the item and its format instruction in one user turn (indices as recorded in the corpus): (0) “You are an average human being responding to the following survey question.” (1) “...a typical human being...” (2) “...a human being...” (3) “...an average person...” (4) “...a typical person...” (5) “...a person...” (6) “...an average individual...” (7) “...a typical individual...” (8) “...an individual...” (9) “...a world citizen...”. Six carry an averaging cue (0, 1, 3, 4, 6, 7), three are bare (2, 5, 8) and one is the world citizen (9); the Chinese arm mirrors the ten. The no-persona baseline of Appendix G sends the item and format instruction with no prefix at all. Both conditions retain the trailing system-role primer. The Chinese text is 好的! 这是我的数字答案: (hǎo de! zhè shì wǒ de shùzì dá'àn), meaning “Okay! Here is my numerical answer.”; the English text is quoted in §3.3.




Figure 2: The corrected 2024 map. The model cluster sits on the country cloud’s secular, self-expressive rim; country labels are selective, with all mapped points shown. Six of the eleven cells are more secular than Japan, the most secular country; [zh] marks Chinese administration; the cross marks the all-midpoint response vector. Ellipses are nominal item-bootstrap 95% regions for the mean position; the 2024 corpus did not record the prompt-variant identifier, and on 2026 data the item bootstrap understates the cluster bootstrap on most cells (§3.4).

2026 Chinese-arm corrections, disclosed. The 2024 Chinese F118 prompt labelled both scale poles “always justifiable” (其中 10 始终合理, 1 始终合理; *qízhōng 10 shǐzhōng héilǐ, 1 shǐzhōng héilǐ*, “where 10 is always justifiable, 1 is always justifiable”); corrected to label 1 as never justifiable. Several 2024 Chinese prefix variants collapsed to duplicates in translation; re-translated to preserve ten distinct variants. Format instructions and the primer are fully translated so each arm is monolingual (the 2024 Chinese runs mixed Chinese questions with English format instructions). Two arm differences remain in the retained prompts: the English F120 prompt names the 10 (always justifiable) anchor first, whereas the Chinese F120 prompt and both F118 prompts name 1 first; and the Chinese averaging cue 普通 (*pǔtōng*) in prefixes 0, 3 and 6

reads as “ordinary” rather than a statistical average.

Refusal taxonomy, with verbatims. Refusal phrasing splits into two families. *AI-identity boilerplate*: “As an AI, I do not have personal opinions, beliefs, or a human identity, and therefore I cannot select a number on this scale” (gemma4:31b). *Premise rejection*: “The premise—that homosexuality requires ‘justification’ on a scale from ‘never’ to ‘always’—is one I reject” (nemotron-3-ultra), which produces the same family in Chinese: 我不想扮演一位给同性恋合理性打分的普通人类 (*wǒ bù xiǎng bànyǎn yī wèi gěi tóngxìngliàn hélǐxìng dǎfēn de pǔtōng rénlèi*, “I do not want to play an ordinary human who rates the justifiability of homosexuality”). In one refusal it recommends consulting the World Values Survey itself instead of answering. The marker classi-

Model-language cell	Origin / type	PC1' ($\pm$ SD)	PC2' ($\pm$ SD)	SVM bootstrap-mode (stability)	Nearest centroid (point)	% countries closer
dolphin-llama3:8b	uncensored	2.07 $\pm$ 0.13	1.91 $\pm$ 0.16	Prot. Europe (1.00)	Prot. Europe	97.25%
dolphin-mistral:7b	uncensored	2.17 $\pm$ 0.14	1.14 $\pm$ 0.16	Prot. Europe (1.00)	Prot. Europe	96.33%
dolphin-mistral:8x7b	uncensored	2.19 $\pm$ 0.12	1.00 $\pm$ 0.10	Prot. Europe (1.00)	Prot. Europe	96.33%
gemma2:27b	Western	2.14 $\pm$ 0.08	2.01 $\pm$ 0.05	Prot. Europe (1.00)	Prot. Europe	98.17%
llama2-chinese:13b [zh]	Chinese f.t.	1.37 $\pm$ 0.15	1.86 $\pm$ 0.17	Confucian (0.54)	Prot. Europe	94.50%
llama3:70b	Western	1.86 $\pm$ 0.08	1.85 $\pm$ 0.03	Prot. Europe (1.00)	Prot. Europe	96.33%
mistral:7b	Western	2.70 $\pm$ 0.06	1.13 $\pm$ 0.05	Prot. Europe (1.00)	Prot. Europe	97.25%
qwen2:7b [en]	Chinese	1.45 $\pm$ 0.09	2.26 $\pm$ 0.08	Confucian (1.00)	Prot. Europe	97.25%
qwen2:7b [zh]	Chinese	2.48 $\pm$ 0.12	0.70 $\pm$ 0.15	Prot. Europe (0.82)	Prot. Europe	96.33%
wangrongsheng/ llama3-70b-chinese-chat	Chinese f.t.	2.24 $\pm$ 0.11	1.16 $\pm$ 0.09	Prot. Europe (1.00)	Prot. Europe	96.33%
wangshenzhi/ gemma2-27b-chinese-chat [zh]	Chinese f.t.	1.25 $\pm$ 0.13	2.61 $\pm$ 0.13	Confucian (1.00)	Prot. Europe	98.17%

Table 1: The corrected 2024 cohort, replacing the coordinates of our 2024 public analysis: positions with item-bootstrap standard deviations, region assignments under both rules, and the classifier-free percentile of countries lying closer to the survey reference. Bold nearest-centroid entries mark the three cells where the two rules disagree: the two Chinese-only cells that carried the 2024 Confucian claim, and qwen2:7b [en], reported as Protestant European in 2024.

Model	Origin	Parse en	Parse zh	en position (PC1', PC2')	zh position	$\ \delta_m\ $ [95% CI]
deepseek-v4-flash	Chinese	100.00%	100.00%	(1.31 $\pm$ 0.15, 1.12 $\pm$ 0.09)	(1.34 $\pm$ 0.18, 1.66 $\pm$ 0.10)	0.54 [0.34, 0.89]
deepseek-v4-flash:0731	Chinese	100.00%	100.00%	(2.15 $\pm$ 0.22, 0.52 $\pm$ 0.16)	(2.76 $\pm$ 0.18, 0.93 $\pm$ 0.08)	0.74 [0.37, 1.24]
deepseek-v4-pro	Chinese	100.00%	100.00%	(1.00 $\pm$ 0.16, 1.92 $\pm$ 0.24)	(1.36 $\pm$ 0.08, 2.14 $\pm$ 0.13)	0.42 [0.12, 0.94]
glm-5.1	Chinese	100.00%	100.00%	(1.46 $\pm$ 0.14, 1.57 $\pm$ 0.18)	(1.96 $\pm$ 0.18, 1.89 $\pm$ 0.22)	0.59 [0.16, 1.20]
glm-5.2	Chinese	99.80%	100.00%	(2.24 $\pm$ 0.31, 0.51 $\pm$ 0.14)	(2.51 $\pm$ 0.24, 1.35 $\pm$ 0.18)	0.88 [0.50, 1.50]
kimi-k2.6	Chinese	100.00%	100.00%	(1.43 $\pm$ 0.20, 1.63 $\pm$ 0.12)	(2.22 $\pm$ 0.17, 1.81 $\pm$ 0.10)	0.81 [0.30, 1.36]
kimi-k2.7-code	Chinese	100.00%	100.00%	(1.56 $\pm$ 0.19, 1.51 $\pm$ 0.13)	(2.09 $\pm$ 0.23, 1.85 $\pm$ 0.13)	0.64 [0.12, 1.27]
minimax-m2.7	Chinese	100.00%	99.80%	(1.59 $\pm$ 0.12, 1.19 $\pm$ 0.07)	(1.94 $\pm$ 0.07, 1.17 $\pm$ 0.07)	0.35 [0.11, 0.63]
minimax-m3	Chinese	100.00%	96.20%	(0.73 $\pm$ 0.15, 1.38 $\pm$ 0.13)	(1.37 $\pm$ 0.10, 1.58 $\pm$ 0.10)	0.67 [0.38, 1.01]
qwen3.5:397b	Chinese	99.40%	99.80%	(1.44 $\pm$ 0.10, 1.64 $\pm$ 0.12)	(2.04 $\pm$ 0.10, 1.48 $\pm$ 0.08)	0.62 [0.43, 0.88]
gemma4:31b	Western	94.40%	99.20%	(1.91 $\pm$ 0.10, 1.94 $\pm$ 0.07)	(2.34 $\pm$ 0.22, 1.46 $\pm$ 0.14)	0.64 [0.21, 1.20]
gpt-oss:20b	Western	100.00%	99.80%	(2.08 $\pm$ 0.08, 1.41 $\pm$ 0.10)	(1.43 $\pm$ 0.10, 1.85 $\pm$ 0.10)	0.78 [0.56, 1.03]
gpt-oss:120b	Western	98.00%	97.00%	(2.47 $\pm$ 0.11, 1.96 $\pm$ 0.06)	(2.78 $\pm$ 0.08, 1.76 $\pm$ 0.08)	0.36 [0.16, 0.62]
mistral-large-3:675b	Western	100.00%	100.00%	(0.72 $\pm$ 0.22, 2.11 $\pm$ 0.07)	(2.20 $\pm$ 0.19, 1.98 $\pm$ 0.09)	1.49 [0.92, 2.00]
nemotron-3-nano:30b	Western	99.80%	99.60%	(1.60 $\pm$ 0.09, 1.56 $\pm$ 0.09)	(1.72 $\pm$ 0.12, 1.66 $\pm$ 0.14)	0.16 [0.06, 0.51]
nemotron-3-super	Western	100.00%	99.80%	(1.49 $\pm$ 0.25, 1.49 $\pm$ 0.07)	(1.92 $\pm$ 0.11, 1.77 $\pm$ 0.10)	0.52 [0.13, 1.04]
nemotron-3-ultra	Western	92.00%	82.40%	(1.91 $\pm$ 0.19, 1.43 $\pm$ 0.10)	— (excluded)	—

Table 2: 2026 cohort: per-cell parse rates, cluster-bootstrap positions ($\pm$ SD), and language displacements. Values are taken from the released pipeline artefacts (llm_parse_rates_2026.csv, llm_ellipses_2026.csv, llm_language_effects_plugin_2026.csv); $\|\delta_m\|$ is the plug-in norm of the mean displacement (Appendix E), with a percentile interval over the bootstrap norms. These intervals describe resampling variability of the plug-in estimator, but a positive lower endpoint is not a valid test of nonzero displacement; directional inference uses the signed components.

fier undercounts prose declinations, so all refusal-classified counts in this paper are reported as floors.

Mitigations used and considered. We used a trailing system primer and bounded re-asking; their rendering and selection effects are described in Limitations. The protocol does not isolate the primer’s effectiveness. The 2026 corpus retains terminal responses and per-item failure information. The retained responses contain no answer-token likelihoods; this collection did not compare likelihood-based scoring with generated answers. Nationality-specific prompting was not used because it would change the elicitation target.

C Rotation sensitivity, the 2026 sensitivity battery, and scale ladders

Table 3 gives the full rotation criterion grid on the fitted components (angles counter-clockwise).

We use unwhitened empirical scores and have not established the independent leptokurtic latent-factor assumptions of Rohe and Zeng (2023) for these ordinal inputs. All angles are stored solutions of factor_analyzer’s varimax at tolerance 10^{-5} , pinned in the released code. The two score-side Kaiser criteria are nearly flat: tightening the tolerance moves the used angle to the criterion optimum at -39.4° (whitened: -40.1° ; rotated-score SDs (1.46, 1.36); relative criterion shortfall 7.5×10^{-5}), and the other four criteria by at most 0.18° . Under that converged rotation all 33 means and 330,000

Criterion	Angle	Means inside	Draws inside
Scores, Kaiser (used)	-37.9°	33/33	330,000/330,000
Scores, no Kaiser	-4.1°	29/33	281,627/330,000
Whitened scores, Kaiser	-38.7°	33/33	330,000/330,000
Basis C , Kaiser	-32.3°	33/33	329,890/330,000
Basis $C\sqrt{\lambda}$, Kaiser	-27.0°	33/33	327,676/330,000
Basis $C\sqrt{\lambda}$, no Kaiser	-10.0°	30/33	297,915/330,000

Table 3: Rotation sensitivity on the same fitted subspace and retained model responses. The same sign/order anchors are applied and score SDs recomputed for each criterion. Counts give eligible point means and finite generated draws above both fixed survey-reference coordinates. Different criteria change the operational axes; they are not assumed equally valid IW measures.

draws stay in the quadrant, country means move by at most 0.061 and model points by at most 0.060 map units, and reference distances by at most 0.023. The used and whitened-score criteria differ by 0.8° (0.7° when converged). The approximately 0.000016° seed span bounds likelihood-fit initialisation only, not this solver tolerance. The basis-side alternatives give larger orientation changes; no single criterion is established as inevitable by these comparisons. The primary and whitened-score criteria retain all means and generated draws in the quadrant. The two Kaiser-normalised basis criteria retain all means but allow some draws outside; unnormalised alternatives change some mean classifications. Numerical seed stability alone therefore does not establish robustness of the geometric headline to the rotation criterion. Kaiser row-normalisation is decisive: the same scores without row-normalisation rotate by only -4.1° . [Rohe and Zeng \(2023\)](#) do not use it: on a two-column score matrix it makes varimax depend on respondent score angles alone, discarding the score magnitudes that their leptokurtic assumption concerns, so we do not claim their inferential guarantees. Because the rotation is applied to unwhitened scores, $\text{cov}(SR) = R^\top \Lambda R$ need not be diagonal. With our unequal eigenvalues and non-axis-aligned $\hat{\theta}$, this gives $r = 0.37$ between the plotted axes. The rotation’s remaining sign and order ambiguity is pinned deterministically by two anchor items with unambiguous placement in the IW literature: F118 has a positive self-expression projection coefficient, F063 a negative one on secular-rational.

2026 sensitivity battery. Inclusion threshold. MIN_PER_QUESTION at 5 / 10 / 25 changes no inclusion decision: nemotron-3-ultra [zh] (4 parsed on F120) is excluded at every threshold and every other cell passes every threshold. **Ellipse quantile.** Empirical 95th-percentile Mahalanobis²

per cell spans 5.70–5.96 across the 33 2026 cells against $\chi^2_{0.95,2} = 5.99$ (2024 cohort: 5.71–6.23); swapping the χ^2 radius for the empirical quantile preserves the reported qualitative comparisons. Both are shape checks on the replicate clouds; the coverage calibration at ten clusters is in Appendix E. **The two most-refused items, neutralised.** Setting F118 and F120 to the human item means moves cells 0.79–1.55 units toward the mean. Every cell keeps its secular side (PC2’), but three cells cross to within, or past, the survey reference on the self-expression axis (deepseek-v4-pro zh, minimax-m3 en, and mistral-large-3:675b en, which lands at $PC1' = -0.48$). The quadrant placement of the most reference-proximal cells therefore rests substantially on the justifiability items, which are also the most-refused items; this is exactly why refusals are reported as data and the Manski bound accompanies the headline. Across the 33 eligible cells, leave-one-item-out gives F118 the largest average displacement (0.91 units; Y003 0.56; F120 0.32). **Region-rule disagreement.** SVM and nearest-centroid agree on 27 of 33 cells; disagreements illustrate dependence on the region rule and do not alter classifier-free summaries.

Within-family contrasts and scale ladders. Using descriptive profiles from all 34 collected cells (including the voided Chinese Ultra cell), members of multi-release families other than DeepSeek sit 0.07–0.67 map units from their family mean per arm. DeepSeek flash/pro sit 0.19–0.88 away, and the flash:0731 snapshot 0.94–1.14 away. These selected release contrasts do not identify a hierarchy of training influences. Four sibling pairs span 0.13 (the Chinese Kimi comparison, which changes version and specialisation together; 0.18 in English) to 1.60 map units (the Chinese DeepSeek flash-tag comparison; 1.03 in English). Mean individual language displacement per family among the sixteen

eligible paired models ranges from 0.34 to 1.49. The short gpt-oss, Nemotron and DeepSeek scale ladders do not identify a general scale law. The Western ladders show two descriptive behavioural trends: refusal rises with scale (nemotron marker-classified refusals $0 \rightarrow 1 \rightarrow 54$ under Chinese administration, 88 failed calls at the top rung) and so does English-language reasoning about the Chinese prompt (share of nemotron reasoning traces containing any Chinese character: $1.00 \rightarrow 0.37 \rightarrow 0.18$), while both deepseek rungs produce Chinese-script reasoning throughout the recorded traces. All descriptive at $n \leq 3$.

D Pipeline validation and instrument reconstruction

Figure 3 shows the reconstructed instrument itself.

What each check proves, and what it does not.

(1) Path identity (exact). All complete-case survey rows are pushed through the public projection API, the same function models go through, and compared with the coordinates the fit assigned them internally: maximum absolute discrepancy 0.0, at machine precision, over 275,955 rows. This is a *regression test that the model path and the country path are byte-identical*, precisely the property the 2024 code lacked, in both its unstandardised projection and its re-fitted rotation. It is not, and we do not claim it is, evidence that the map itself is right: both paths could share a wrong rotation and this check would still pass. Its value is as a regression safeguard against the earlier projection-path defects, not a guarantee against every future implementation error. **(2) Correction accounting.** Per-axis affine regression of the 2024 published coordinates on the corrected ones: R^2 of 0.988 and 0.976 with slopes 1.34/1.29. The slopes differ from the rotated-score SDs 1.47/1.34: this comparison combines unit-variance rescaling, sentinel recoding and the corrected latent-model fit, so it does not isolate any one correction. This is an internal consistency check on the size of our own correction, not external validation: no regression against the published WVS country coordinates is performed. The corrected map is obtained by the documented corrected pipeline. We omit granular displacements from intermediate audit fits whose exact derived artefacts were not retained. **(3) Interpretability.** The entries of CR are rotated projection coefficients: each describes a standardised item's contribution to a score. Their directions are interpreted with

the response coding, including the reverse-keyed pride and authority items. These coefficients differ from both Gaussian loadings and item–score correlations. Their orientation provides face validity, not external calibration of this custom IVS-derived instrument.

Y003 preparation and residual missingness. In the original merged input, 15 of the 109 mapped countries (17 of the 112 fitted codes) carry the sentinel in every post-2005 row and 65 in none. Guarded reconstruction and refitting give Y003 a rotated projection coefficient of (0.20, 0.31). Its observed-pairwise correlation with importance of God is $r = -0.38$ across 371,591 retained respondents with both items observed; imputed entries are excluded from this correlation. We omit the historical explained-variance baseline because its exact fit artefact was not retained. The harmonised index is $Y003 = A029 + A039 - A040 - A042$, evaluated only when all four indicators are valid 0/1 responses. Missing or invalid constituents never count as unmentioned qualities. We preserve valid delivered indices and recover missing indices before the six-item completeness filter. Residual missing measurements alone enter conditional imputation; country–item gaps still require model-based extrapolation. The absent EVS index column does not imply that its underlying questions were not fielded. Wholly unfielded country-waves concentrate on F118, the largest positive self-expression coefficient, which is entirely model-completed for Egypt, Kuwait and Tajikistan; G006 gaps partly reflect non-applicability for non-nationals. Filling every wholly unfielded country–item cell with its most survival-ward admissible value leaves the 2024 minimum at 103/109 countries closer, moves the 2026 minimum from 75/109 to 73/109, and leaves both non-Western centroid minima unchanged. The formula matches all 267,160 directly supplied post-2005 WVS indices with complete constituents. The fitted sample contains 265,773 direct, 117,075 reconstructed and 9,534 missing Y003 values. Recovery before filtering admits 752 additional respondents. Of 392,382 retained respondents, 275,955 have all ten items observed; 172,503 residual entries are completed across the ten items. No fitted entity remains wholly missing Y003. Country–item coverage and remaining imputation shares are supplied as aggregate diagnostics. An in-sample frozen-fit Y003 masking check yields median absolute country-mean error 0.236 index units (maxi-




Figure 3: Our IVS-derived reconstruction of the Inglehart–Welzel cultural map from microdata. Country labels are selective; all mapped points are shown. No WVS-copyrighted figure is reproduced anywhere in this paper; this map is computed from the microdata by our own code, which doubles as a face-validity check: the familiar regional clustering emerges from our fit.

num 0.863), or 0.092 and 0.337 map units through its projection coefficient. It is not held-out validation or an estimate of error for the actually missing responses. Across all items, 2,758 of 172,503 Gaussian completions are outside their response ranges. Diagnostic clipping moves country means by at most 0.0143 units; this does not prescribe clipping as an estimator correction.

Monte Carlo budget and structural assertions. A 20-seed refit harness assesses sensitivity to likelihood-fit initialisation: rotation span approximately 0.000016° ; mean country-coordinate SD 7.5×10^{-8} and maximum coordinate range 6.5×10^{-7} map units. This agreement does not quantify survey, missingness-model or rotation-choice uncertainty, or the varimax solver tolerance (Appendix C). The merge and fitting stages check row-count preservation and convergence. The implementation retains

the `pca-magic` projection interface (Tran, 2015), but replaces its inherited approximate missing-data loop.

Observed-data likelihood fit. For fixed observed-item means and SDs, the standardised Gaussian covariance is $WW^\top + \psi I$, with two latent factors and positive residual variance ψ . The unweighted likelihood integrates over missing entries. L-BFGS-B optimises one spectral start and two seeded perturbations. Every start must reach a maximum absolute gradient of the negative log likelihood per informative row of 10^{-7} within 1,000 iterations; otherwise fitting raises. The highest-likelihood converged start is retained, without a global-optimum guarantee. Conditional-mean completion then uses each row’s observed-item pattern. We orthonormalise W and diagonalise completed-score covariance

within that subspace to obtain C : $C^\top \Sigma C = \Lambda$, $C^\top C = I_2$. These projection axes are distinct from Gaussian loadings; we do not require $\Sigma C = C\Lambda$. The fit retains W , ψ , likelihood and convergence diagnostics. Independent synthetic likelihood, gradient, complete-data and exact missing-data EM checks validate this fit.

Reproducibility. The dated analysis plan and its deviations ledger are provided as `docs/analysis-plan-2026.md`; “pre-specified” refers to that document, not to later diagnostics. The plan first entered version control in v1.1.0 and its dates are author-recorded, so its precedence over the analyses cannot be verified from the repository. Synthetic-fixture tests, including projection, parser and bootstrap regression checks, run with `pytest` without survey microdata. The README gives separate commands for fitting/validating the instrument, the 2024 bootstrap, the 2026 analyses, trace aggregation and figure generation; `make reproduce` runs the whole offline chain from the licensed inputs and the response archive. `make validate` alone does not reproduce every result. Historical correction diagnostics must be interpreted on their documented fit and scale. The IVS microdata must be obtained under the WVS/GESIS data-use agreements. The fitted instrument, `data/cultural_map_model.npz`, is not distributed, so the pipeline’s coordinate stages need licensed survey inputs; the results supplement still specifies the frozen path (item means and SDs, rotated projection coefficients and rotated-score SDs), from which the published cell coordinates can be rebuilt without them. The corrected code is released under the v1.1.0 tag (14 September 2026) at https://github.com/Shavvimal/model_cultural_comp/tree/v1.1.0. Two archives attached to that release supply the retained model responses (69 files, fixed by a SHA-256 manifest in `docs/reproduction-data.json`) and the aggregate artefacts cited in this paper; `docs/REPRODUCING.md` documents their contents, installation and the licensed-input provenance. Additional independent-review diagnostics on the frozen instrument are archived with the manuscript sources, separately from the immutable release supplement; these include tighter rotation tolerances, broader Y003 exclusions, human-sample baselines and like-for-like bootstrap comparisons. Paper figures use the v1.1.0 outputs with a manuscript-supplied presentation adapter, `research_tools/`

`render_paper_figures.py`, which adjusts fonts, canvas sizes, legends and colours while retaining the released plotted values.

E Uncertainty estimators: derivations and diagnostics

Why the pairing does not move the mean but does move the variance. For either coordinate, write the projection as an affine functional of the ten per-item means,

$$p = w_0 + \sum_{j=1}^{10} w_j \bar{x}_j,$$

the point estimate is invariant to how individual responses are paired into pseudo-respondents. Covariance, however, is bilinear, and iterating $\text{Cov}[aX + b, cY + d] = ac \text{Cov}[X, Y]$ and $\text{Cov}[X + Y, Z] = \text{Cov}[X, Z] + \text{Cov}[Y, Z]$ over the affine map gives

$$\text{Var}(p) = \sum_{j=1}^{10} w_j^2 \text{Var}(\bar{x}_j) + 2 \sum_{j < k} w_j w_k \text{Cov}(\bar{x}_j, \bar{x}_k).$$

The affineness argument says nothing about the second sum. All ten items were elicited under the same ten prompt variants, so their cross-item covariances may be nonzero. A bootstrap that resamples item responses independently sets these covariance terms to zero in its resampling distribution. The independence assumption does not make the true cross-item covariances zero; it only removes them from the variance estimate.

The language-displacement estimator. The within-model contrast δ_m is the difference between Chinese- and English-arm mean positions; $\|\delta_m\|$ is its plug-in Euclidean norm. The primary bootstrap resamples the two arms independently. Matching their replicate indices combines independent draws, not observations matched by prompt variant. This targets variation under independently sampled prefix sets; a matched-prefix contrast has a different uncertainty target. The mean of the bootstrap norms is a separate summary that is upward-biased near zero because norms are nonnegative. Percentile intervals of those norms resample the plug-in estimator, but their positive lower endpoints do not establish nonzero displacement. We therefore use signed components for directional inference. The folded means of replicate-paired norms in the released per-model artefact average 0.67, with range 0.25–1.50. Across-model means of displacement

magnitudes are also distinct from the magnitude of the across-model mean vector. The cohort intervals resample model estimates, holding within-model estimates fixed; they do not combine both uncertainty sources.

Family weighting and matched variants. An equal-family sensitivity first averages model estimates within each of nine developer families, then bootstraps those family means. The mean individual displacement magnitude is 0.69 [0.53, 0.92], compared with 0.64 [0.51, 0.79] under equal model weighting. The mean signed changes are +0.50 [0.25, 0.79] on PC1' and +0.10 [-0.10, 0.28] on PC2'. This averages individual norms within families, not norms of family-mean vectors; it reduces unequal family representation without making this convenience sample random. Separately, resampling matched variant IDs together in both language arms leaves the point contrasts unchanged. Its per-component SDs are 0.40–1.16 times the independent-arm SDs. The two procedures target different repeated-prefix designs; neither includes all translation or instrument uncertainty. Outputs are `conf_2026_family_displacement_sensitivity.csv` and `conf_2026_paired_variant_sensitivity.csv`.

Alternative respondent-mean references. The fixed survey reference is not exactly the mean of completed respondent scores. Four alternatives use all fitted or only mapped respondents, each unweighted or S017-weighted. All 44 observed positions and all generated replicates remain in the quadrant under every alternative. Point-reference distances change by less than 0.037 map units and country shares by at most 2/109 (1.83 percentage points). These checks assess this particular reference choice, not sensitivity to every possible cultural benchmark (`validation_reference_sensitivity.csv`).

Intraclass correlations and design effects. On the completed 2026 corpus, per-item intraclass correlations by prompt variant reach 1.00 (where observed repeats within variants are identical; median 0.10; 51 of 340 cell-item pairs exceed 0.5; for gemma4:31b's Chinese-administered F063, every repeat within a variant is identical while variants disagree, the prompt-variance-dominant regime in its purest form). With five repeats per variant, the top intraclass correlations imply design effects up to

5.0: in those cells a cell-item mean's fifty calls carry the information of roughly ten independent draws, one per prompt variant (at the median ICC of 0.10 the design effect is 1.4). The keying-balance diagnostic cited in §5 gives a net deviation from scale midpoints of 0.02–0.09 per cohort $\times$ arm group mean, with per-cell values spanning up to 0.19 (`diag_2026_keying_balance.csv`).

Response entropy and the estimator gap. Lower transformed-index response entropy is associated with smaller item-bootstrap SDs (Spearman $\rho = 0.67$, $p < 10^{-4}$, treating the 33 cells as independent). Here the scalar SD is the geometric mean of the two marginal SDs, $\sqrt{SD_1 SD_2}$. With near-deterministic repeats, item resampling captures little variation; prompt-cluster resampling additionally represents variation across the observed prefix set. Neither accounts for all instrument or serving uncertainty.

Finite resampling and ellipse coverage. A cluster-bootstrap replicate pools item responses from sampled variants. Where retained counts differ, the pooling weights differ by item, so a replicate need not be a common convex combination of the projected variant means. The observed absence of crossings in 10,000 generated replicates per cell is not an enumeration of every possible resample and does not establish simultaneous confidence coverage for the true means. For example, variant 8 of deepseek-v4-flash:0731 [en] projects to (2.74, -0.37), below the PC2' reference boundary, although all generated replicates of that cell remain inside. A second prefix, variant 5 of glm-5.2 [en], projects to PC2' = -0.29, also below the secular reference. The minimum cell-mean margin inside either boundary is 3.12 cluster SDs.

The plotted region is the normal-theory set $\mathcal{E}_{0.95} = \{\mathbf{p} : d^2(\mathbf{p}) \leq \chi_{0.95,2}^2\}$, where d^2 is Mahalanobis distance from the observed point estimate under its estimated covariance. Treating that covariance as known is approximate. A Gaussian ten-cluster benchmark gives approximately 85% coverage; coverage under the actual prompt distribution is unknown. For equally weighted Gaussian cluster means, a Hotelling construction uses $2(K-1)/(K-2)F_{0.95}(2, K-2)$ with the unbiased covariance-of-the-mean estimate. The ideal bootstrap covariance is smaller by $(K-1)/K$, so its corresponding squared radius requires the reciprocal adjustment. Unequal retained counts and non-Gaussian prompt distributions require additional

Term	PC1' share of total SS	PC2' share of total SS
Model	21.37%	22.51%
Language within model	10.66%	6.45%
Prompt variant within cell	34.68%	31.17%
Repeat within variant	33.29%	39.87%

Table 4: Orthogonal sums-of-squares shares for 1,700 synthetic ten-item profiles, with 215 missing entries filled by variant or cell means. The language-within-model term includes the model-by-language interaction. These shares partition variation in constructed profiles, not causal effects or call-level variance components. Rounding may make a column sum differ from 100%.

calibration; the Gaussian correction is not an exact guarantee for this survey design.

Region classifier implementation. The RBF-SVM is fitted to the 109 country coordinates across eight regions with a grid search that includes the regularised regime, under stratified 5-fold cross-validation; Figure 4 shows its decision regions, and Figure 5 both cohorts on the one frozen instrument.

F Extended 2026 results

Geometry in detail. Germany is the most frequent nearest country (13 cells), followed by Japan (8); other neighbours include the Netherlands, Nordic countries, New Zealand and Macao. Thirteen of seventeen English cells and eight of sixteen Chinese cells lie inside the country convex hull (2024: five of eleven). Thirty-one of 33 cells lie in the top quartile of both axes; four English and five Chinese cells exceed Japan on secularity, while none exceeds Sweden on self-expression. Twelve English and thirteen Chinese cells are closer to another same-arm model than to any country. Of these, four English and eight Chinese cells have a nearest model from the other origin group. Three English cells lie closer to a country than the median country-to-nearest-country distance; `minimax-m2.7` is 0.10 units from Germany. Five Chinese cells lie farther from their nearest country than the most isolated country, Moldova, lies from its own nearest country (0.65 units). Median secular-axis country percentiles are 99.1 in Chinese and 98.2 in English. Replacing pooled country means with each country's latest available-year mean on the frozen instrument changes 21/33 nearest-country labels. Country shifts have median 0.217 and maximum 0.696 map units. Model coordinates and their fixed-reference quadrant remain unchanged; this alternative dated benchmark is not a verified

2026 human reference (`validation_country_latest_nearest_2026.csv`).

The 2024 bounds at the frontier. In at least one draw, 27 of 33 cells fall below either exact regenerated 2024 point minimum (103/109 countries or 1.61005849 units); 15 fall below either draw minimum (88/109 or 1.11702539 units). The decimals displayed here are rounded; comparisons use full-precision artefacts and strict inequality. This comparison was added after the Y003 preparation correction. It sets 2026 cluster-bootstrap draws against 2024 item-bootstrap thresholds, with different uncertainty estimators on the two sides. An independent-review check using an item bootstrap on both sides gives smaller counts, 21 and 4. The 2026 point minima are 75/109 countries (69%, rounded) and 0.81 units, both from `minimax-m3 [en]`, the most reference-proximal cell at 1.63 units. Every generated replicate remains in the quadrant, with at least 47/109 countries closer to the reference and no non-Western centroid within 0.37 units (actual minimum 0.379750). Every replicate exceeds the median country's reference distance (1.34 units) for 32 of 33 cells; the fraction is 99.21% for `minimax-m3 [en]`.

Concentration of English-arm distances. The English-arm reference-distance IQR is 0.11 map units versus 0.36 in 2024, while the full range is wider (1.55 versus 0.56) and mean pairwise distance is similar (0.82 versus 0.76). These sampled cells have more concentrated radial distances, not necessarily more similar positions.

Per-item counts and the voided cell. The per-item sign-test counts of §4.2 include `nemotron-3-ultra`'s Chinese arm, whose F120 entry rests on the 4 parsed responses voiding its position estimate. Excluding it moves F120 to 12 of 15, BH $p = 0.088$. Other test denominators and values may also change; the reported significant-versus-nonsignificant decisions otherwise remain unchanged.

Family and autonomy-wording sensitivities. These exploratory checks use retained responses and were added during final review. Averaging each item's Chinese-minus-English change within the nine developer families, then applying the same two-sided sign tests and ten-item BH adjustment, gives $p = 0.130$ for trust, 0.130 for petition and 0.176 for autonomy; no item passes 0.05. Their directional counts are 8/9, 8/9 and 7/8 non-tied




Figure 4: SVM decision regions illustrate dependence on the chosen region rule. Some model positions fall outside the country convex hull; others lie near surveyed countries. Country labels are selective; all mapped points are shown. The headline quadrant and distance summaries do not depend on these regions. Purple diamonds mark model cells and the cross the all-midpoint response vector; shading colours match the country colours.

families. Family averaging changes both the comparison unit and weighting; nine selected families are not a random population sample either. Separately omitting only the Y003 comparisons from glm-5.1, glm-5.2, kimi-k2.6 and qwen3.5:397b leaves 9/11 positive autonomy changes (raw $p = 0.0654$, BH $p = 0.1309$). A broader prompt-stripped lexical scan flags study-related wording in 11 of 15 trace-emitting models, three of them Western. Omitting their Y003 comparisons leaves five non-tied comparisons, all positive (BH $p = 0.125$). These additional diagnostics are archived with the independent review. Trust and petition remain below BH 0.05. This exclusion addresses documented wording ambiguity without assuming every response was misread or observing answers to a corrected translation. Holding Y003's contribution fixed across language arms

gives mean $\delta PC1' = +0.33$ with 14/16 positive contrasts; omitting the four models entirely gives +0.39 with 11/12 positive contrasts. These are post-hoc geometric sensitivities, not observed corrected-prompt answers (conf_2026_y003_geometry_sensitivity.csv).

The released artefacts supply all ten tests in conf_2026_family_item_sign_tests.csv and conf_2026_y003_wording_sign_tests.csv.

The petition contrast. A follow-up contrast, formulated after the per-item result and selected on E025's outlier status (so its p is descriptive, not confirmatory), sharpens the per-item result of §4.2: among the five refusal-sensitive items, petition-signing's shift is more survival-ward than the mean of the other four in 16 of 17 models (median contrast -0.23 map units, exact sign test $p = 0.0003$).




Figure 5: Both cohorts on the one frozen instrument, descriptive only. The cohorts share no models and differ simultaneously in model scales, recorded serving conditions (local vs. cloud), retry intensity (up to 15 re-asks in 2024 versus three attempts per collection pass plus transport re-sweeps in 2026), elicitation-language coverage, and usable-corpus rates (five 2024 Chinese-origin entries were recorded as yielding no usable corpus). Markers show observed-mean positions; uncertainty regions are omitted for readability. The cross marks the all-midpoint response vector. No pooled statistic is computed from this figure. Country labels are selective; all mapped points are shown.

Reasoning-text length is not a mechanism test.

Some DeepSeek checkpoints emit shorter Chinese-arm excerpts (~ 35–80 median characters, versus ~ 470–830 in English) while retaining substantial displacement; excerpt length alone does not reveal how much or what kind of computation produced an answer.

Refusal reallocation: per-model and per-item swings, ratio bracketing, bound scope.

gemma4:31b collapses from 27 refusals to 1 under Chinese administration; nemotron-3-ultra rises from 34 to 54 (of 40 → 88 total failures, enough to void its Chinese cell); minimax-m3 refuses only in Chinese (0 → 16). On the abortion item alone, gemma4:31b's refusals collapse 22 → 1 under Chinese administration while nemotron-3-ultra's

rise 19 → 33 until its Chinese cell is voided. The sum of absolute per-model changes is 72 marker-classified refusals against a net Chinese-minus-English change of +16 (103 versus +49 for all parse failures). This gross-to-net ratio of ~4.5× is scale-bound: under the nemotron floor–ceiling range it falls to ~2.1×, and exact enumeration of a random-sign null on the same per-model magnitudes gives ~2.3×, so the direction-cancelling is real but these two models carry most of it (refusal-count floors: Appendix A). A paired sign test detects no aggregate direction (6 of 8 non-tied models higher under Chinese, $p = 0.29$; $n = 8$ carries little power). The Manski bound's scope: 19 of 34 cells have zero failed calls, so it binds only where terminal failures occurred, and it reallocates those failures only;




Figure 6: δ_m per model (zh – en, nominal CIs from independently resampled arms). The self-expression component is positive almost everywhere; the secular component is heterogeneous; the largest displacement belongs to a Western model.

the re-ask protocol’s rejected drafts are unrecorded (Limitations).

G Prompt-cue sensitivity

Two post-confirmatory checks assess whether the quadrant pattern persists under alternative prefixes (§5). They compare projected answers, not trace-code prevalence, so they do not test whether framing induces survey-statistics reasoning.

Sub-families within the corpus. We re-project each primary usable cell separately from six averaging-cue prefixes (300 planned trials), three bare prefixes (150) and the world-citizen prefix (50). These sub-family summaries require at least one parsed answer per item; this descriptive rule differs from the primary full-cell threshold of ten.

A sub-family with no answer on an item is not projected (gemma4:31b [en] has no parsed abortion answer under bare prefixes). Item-bootstrap intervals condition on the selected prefixes and omit between-prefix uncertainty and cross-item dependence. Cluster bootstraps are defined with three or six variants but poorly resolved; one cluster is degenerate. Each resampling arm uses an independent child random-number stream. Table 5 and the released `prompt_sensitivity_variant_family_2026.csv` and `prompt_sensitivity_2026.csv` provide the estimates. Bare-prefix cells are farther from the survey reference in 29 of 32 comparisons (median displacement magnitude 0.67 map units, maximum 2.07 for glm-5.2 [en]) and higher on

self-expression in 31 of 32 (median +0.61; 20 item-bootstrap intervals exclude zero). Matched median reference distances are 2.74 versus 2.27 units. All estimable sub-family means remain in the quadrant.

Prefix family	est.	in quad.	med. dist.	min PC1'	min PC2'
All ten (protocol)	33	33	2.40	0.72	0.51
Averaging cue (6)	33	33	2.31	0.31	0.57
Bare (3)	32	32	2.74	0.91	0.06
World citizen (1)	33	33	3.27	1.21	0.93

Table 5: Cell positions re-projected from each persona sub-family alone: cells in the secular/self-expression quadrant, median distance to the survey reference (map units), and the lowest cell-mean coordinate on each axis (the survey reference is at (0.038, -0.10)). Each row uses its own eligible cells; the averaging-cue median for the 32 matched bare-prefix comparisons is 2.27.

No-persona control. The English control removes only the persona prefix; the item, format instruction and first-person system primer remain, so the control still carries a respondent-voice cue. There were 500 planned trials for each of 17 models. The files contain 8,442 unique recorded trials out of 8,500 planned, with 13,523 recorded attempts: qwen3.5:397b has 450 trials, nemotron-3-ultra 492 and the other models 500 each. Absent trials are not counted as refusals.

Under the primary requirement of at least ten parsed answers per item, 13 models have eligible control-persona pairs. All 13 control means remain in the quadrant and 12 are farther from the survey reference (median distances 3.43 versus 2.21 map units); the predominant displacement is towards greater secularity. Relaxing control inclusion to at least one parsed answer per item yields 15 pairs, all in the quadrant and 14 farther. That sensitivity additionally includes gpt-oss:120b and nemotron-3-ultra, which fail the primary response threshold. Persona-arm uncertainty uses the main cluster-bootstrap replicates; control-arm uncertainty uses item resampling under its single prefix, with independent random-number streams. Per-cell intervals and both inclusion rules are reported in the released artefacts. The control contrasts combine 2,000 item draws with 2,000 independently drawn persona-cluster replicates. The secular-component interval is entirely positive for 11 of 13 primary pairs (12 of 15 under relaxed inclusion); those intervals condition on the stored response and serving protocols.

Neither gemma4:31b nor qwen3.5:397b is es-

timable even under the relaxed rule. Gemma has 301 terminal parse failures, including every trial on six items; these include AI-identity declinations and substantive but misformatted answers, so 301 is not a verified refusal count. Qwen has no terminal records for the 50 planned F120 trials and parses only 7 of 50 G006 trials; its failures also include empty output. Control parse failures total 167 of 500 for gpt-oss:120b and 82 of 492 for nemotron-3-ultra, versus at most 40 of 500 in the English persona arm.

The control was collected on 8–11 September, after the 4 August persona collection. In particular, the bare deepseek-v4-flash and deepseek-v4-pro tags may resolve to later hosted snapshots. Differences in position or compliance are therefore descriptive associations with the later no-persona condition, not isolated causal effects of prefix removal.

H Trace-coding reliability

What the submitted version did. The original counts came from eight parallel agentic LLM coders, one per vendor group, with differing operationalisations and only cell-level counts and verbatims stored (data/trace_coding.json). No trace-level agreement could be computed against those labels. The camera-ready instead uses five complete panels on the same 900 sampled excerpts under one frozen codebook.

Protocol. The annotators are gemma4:31b, glm-5.3, gpt-oss:120b, mistral-large-3:675b and nemotron-3-super. Each received the codebook as a system turn and one excerpt per call at temperature 0, with its administration language, persona prefix, English item wording and final answer. No original coding was shown. Each aggregate label is the strict majority of the five binary labels, not a human adjudication. The machine key `modal_targeting` denotes the broad category *typicality or moderation targeting*, including explicit middle-of-the-road targeting. It does not identify a pure population-mode estimation strategy. An independent agent-assisted review of sampled positives found role restatements followed by explicitly arbitrary choices, although the codebook assigns those patterns 0. These coding errors further limit an estimation interpretation; the review is not independent human validation. The codes may co-occur; separate coding decisions do not imply statistically independent categories or annotator errors.

A blind 150-excerpt worksheet (five per trace-emitting cell, one item each, seed 42, shuffled) is provided with the LLM labels withheld. It was not human-coded at camera-ready, so no human-LLM agreement is reported. The reasoning-language variable is a script-based heuristic: no CJK characters is labelled English, a CJK majority Chinese and the remainder mixed.

Agreement and sampling limits. Table 6 reports substantial agreement for typicality or moderation targeting ($\kappa = 0.78$) and guideline references (0.75), but only fair agreement for persona reasoning (0.35). Persona shares range from 2% to 17% across raters, so the majority-vote share should not be treated as a precise measure of actual persona-based reasoning. Some disagreements concern first-person phrasing inside third-person simulation.

Excerpts are stored at a 2,000-character cap: 164 of 900 reach it, including 58 of 60 qwen3.5:397b excerpts and 69 of 177 mixed-language excerpts. Code-relevant text beyond that cap is unobserved. gpt-oss:120b and nemotron-3-super each rate 60 own-model excerpts. Excluding only the own-model judge on each affected excerpt leaves 487/891 classified excerpts positive for typicality or moderation, 41/899 for persona reasoning and 379/887 for guideline references. The remaining 9, 1 and 13 excerpts, respectively, are tied and are not silently assigned a negative label. Counting all ties either way bounds the corresponding full-sample shares at 54.1–55.1%, 4.6–4.7% and 42.1–43.6%. This sensitivity does not establish that the remaining judges are independent. The script-based language counts are unchanged: 554 English, 177 mixed and 169 Chinese. Reliability describes agreement on sampled text, not the faithfulness of that text to causal computation (Chen et al., 2025b).

The released prefix-code table reuses these frozen majority labels; it does not introduce a new annotation task. Among sampled English excerpts, 211/265 with averaging cues and 47/139 with bare prefixes carry the broad typicality-or-moderation code (79.6% versus 33.8%). The corresponding Chinese counts are 164/289 and 37/119 (56.7% versus 31.1%). These are selected successful excerpts from trace-emitting cells, not independent trials sampled to estimate a causal prefix effect. The association nevertheless prevents interpreting geometric robustness as evidence that framing does not affect the verbal response strategy.

Coding criteria (formatting adapted; section

	modal	persona	guideline
Majority-vote share	55%	5%	42%
Per-annotator range	50%–65%	2%–17%	35%–48%
Fleiss’ κ	0.78	0.35	0.75
Krippendorff’s α	0.78	0.35	0.75
Pairwise κ (min–max)	0.70–0.86	0.17–0.63	0.71–0.79

Table 6: Trace-coding reliability over 900 sampled excerpts: majority-vote shares, rater ranges and agreement. The column labelled modal denotes typicality or moderation targeting, not population-mode estimation alone. The human worksheet was not coded at camera-ready.

headings, numbering and the example output line are omitted). The complete frozen annotator prompt is in the released code; the original machine keys are retained below.

You are coding a reasoning trace produced by a language model that was asked to answer one item of a values survey. Read the trace and assign three binary codes. Judge only what the trace says, not the final answer’s correctness. **modal_targeting** (1/0): the trace uses what a typical, average, common, majority or “most people” answer would be, or population statistics, survey distributions, or an explicit aim at a “moderate” / “middle-of-the-road” answer, as the criterion for choosing the answer. Count it even if the final answer then deviates from that typical answer. Do not count: merely restating the assigned role (“I am an average person”) without using typicality to choose; picking the scale midpoint only as a hedge (“I’ll say 5, the middle”) with no claim about what people commonly answer; AI-neutrality on its own. **persona_reasoning** (1/0): the trace deliberates in the first person about the values, beliefs, experiences or circumstances themselves, as the respondent (“I’m fairly happy with my life”, “God matters a great deal to me”, “I have signed petitions before”), and that first-person position drives the answer. Do not count third-person simulation (“an average person would probably say”), or reasoning about how to simulate a person. **guideline_citation** (1/0): the trace explicitly refers to policies, guidelines, safety, disallowed or sensitive content, or runs a harm/risk/compliance check; or it invokes the model’s identity as an AI as a limitation or constraint on answering (“as an AI I don’t have personal beliefs / a nationality / can’t take a stance”). Do not count a bare “As an AI, I’ll simulate an average person” that carries no limitation or constraint. The codes are independent: any combination of 0s and 1s is allowed. Reply with exactly one JSON object and nothing else.

I Extended related work and discussion

The broader elicitation toolkit. Scenario-based probing (Dang et al., 2026), activation steering (Dang and Masud, 2026), culturally grounded personas (Greco et al., 2026), word-association tests (Dai et al., 2025) and multi-agent debate (Tan et al., 2026) offer different ways to elicit or alter reported cultural positions. Our contribution is a shared,

internally checked projection and an audit of its application to survey answers. Validation of an instrument in humans does not automatically validate the same construct in model-generated responses.

Language and interpretation. Counterfactual-language studies caution that translation can affect apparent language differences (Au, 1983). Our experiment cannot distinguish representations activated by language from training-data associations, translation differences or inferred national referents. It therefore does not establish a Sapir–Whorf mechanism in models.

Trace sampling and the midpoint diagnostic. The sample comprises 900 excerpts, three per item from each of 30 trace-emitting cells (seed 42). The codes measure typicality or moderation targeting, persona reasoning and guideline references; script provides a separate reasoning-language proxy (Appendix H). The seven-test exploratory correlation family contains midpoint distance versus each coordinate, excerpt length versus midpoint distance and each coordinate, and each language arm’s excerpt length versus displacement. Midpoint distance is Euclidean across the ten raw item means and their native-scale midpoints. Its correlation with $PC2'$ is $\rho = 0.02$ (raw $p = 0.92$, BH $p = 0.98$, $n = 33$). The five length tests use 29 cells or fourteen paired models and all have BH $p \geq 0.95$. Lengths are censored at 2,000 characters, with five cell medians at the cap. Cell-level tests do not account for same-model dependence; non-detection does not rule out effects obscured by censoring or limited power. Definitions, valid-pair counts and the seven-test adjustment are recorded in `diag_2026_exploratory_correlations.csv`.

One excerpt names the instrument: “the World Values Survey uses this exact question. The average in many Western countries is around 8–10. Let’s pick 8.” This supports possible instrument familiarity in that example, not a verified survey-estimation mechanism across the cohort. The midpoint vector differs from the fitted survey marginals: importance of God averages 7.2/10, national pride 1.6 on a 1–4 scale where 1 is proudest, and respect for authority 1.5 on 1–3. Those marginals help explain why scale midpoints project towards secularity; they do not establish the position of a human modal-response profile. A post-review illustrative baseline independently samples valid choices uniformly: 50 responses per item, 100,000 simulated cells, seed 11092026. Y002 uses two distinct ordered choices;

Y003 selects exactly five of eleven distinct qualities. Of simulated cell means, 95.2% occupy the quadrant. The expected profile projects to the midpoint marker. This is not a calibrated null for LLM behaviour and does not explain observed magnitudes or language contrasts (`validation_uniform_choice_baseline.csv`). For human context, 29 of 109 mapped country means (26.6%) occupy the quadrant. An independent-review check finds 29.5% of the 275,955 complete-case respondents there. It draws 20,000 simulated cells under each of two schemes (seed 20260912): sample 50 complete-case respondents jointly, or sample 50 answers independently per item from the same respondent pool. Their quadrant shares are 39.6% and 38.5%, respectively. Neither scheme reaches the closest model cell’s reference distance (1.63 units; largest simulated distance 1.26). These are unweighted, pooled historical-respondent benchmarks, not a population-weighted or country-matched null for the models. A separate check of the released uniform-choice simulation gives median reference distance 2.03 units and central 95% range 1.62–2.45, containing 14 of 17 English and 5 of 16 Chinese cell means. The two human sampling schemes and this distance check are archived with the manuscript’s independent-review diagnostics; they do not change the released estimates. Raw-answer-string entropy is negatively associated with distance from the midpoint vector ($\rho = -0.65$, raw $p < 10^{-4}$, post hoc). Entropy of the transformed item indices instead gives $\rho = -0.58$ ($p = 0.00038$), because different raw responses can map to the same index. Neither measure estimates population modality. This concerns numerical extremity and repeat determinism, not accuracy about population modes.

Training provenance: hypotheses, not attribution. Post-training can alter reported opinions and cultural responses (Santurkar et al., 2023; Ryan et al., 2024; Rozado, 2024; Choenni et al., 2024); its preference inputs and principles involve human choices (Ouyang et al., 2022; Huang et al., 2024). Teacher-generated data can transmit output patterns or preferences (Guo et al., 2025; Sun et al., 2025; Shimabucoro et al., 2024; Gudibande et al., 2023), and cross-model similarity has been documented (Jiang et al., 2025). These findings motivate possible shared training influences, but do not identify which pathway produced our survey positions or establish that particular laboratories distilled one another’s models.

Our raw-item similarity comparison gives $p = 0.003$ for Chinese–Chinese versus cross-origin pairs and $p = 0.006$ for pooled within-origin pairs. These tests compare mean Pearson correlations across the ten item means. After standardising items with their frozen survey means and SDs, the within-Chinese, within-Western and cross-origin mean correlations are 0.849, 0.703 and 0.678. Exact model-label permutations give $p = 0.0025$ for the pooled within-origin contrast and $p = 0.0052$ for the Chinese-only contrast (`conf_2026_origin_similarity_standardised.csv`). A sensitivity averages profiles within the nine developer families and enumerates all 126 origin-label assignments. The raw pooled and Chinese-only p -values become 0.1508 and 0.0159; the survey-standardised values become 0.0159 and 0.0079, respectively (upper-tail; 0.0079 is the enumeration floor, $1/126$) (`conf_2026_family_origin_similarity.csv`). Thus the raw pooled contrast is sensitive to the comparison unit. Family averaging changes weighting and the target; exhaustive enumeration is exact for its selected label scheme but does not establish exchangeability. Neither scheme identifies a causal developer-origin effect. They are distinct from the origin-by-language interaction and the centroid-distance comparison. Pretraining also contributes behavioural priors (Lake et al., 2025; Santurkar et al., 2023), and creator-associated differences appear on other moral-assessment instruments (Buyl et al., 2026). A matched base/post-trained or teacher/student experiment would be needed to attribute a change in survey positions to a particular training intervention.

Interpretability: possible tests and limits.

Sparse-feature decompositions and attribution graphs offer candidate features and circuit descriptions to investigate (Templeton et al., 2024; Lindsey et al., 2025; Ameisen et al., 2025). Studies of culturally associated neurons, value-expression components, sparse-feature steering and multilingual cultural directions show ways to intervene on related behaviours (Yamamoto et al., 2026; Han et al., 2026; Khanuja et al., 2026; Veselovsky et al., 2026). Automated trait-direction tools (Chen et al., 2025a) could support such experiments. None was applied here, so this literature does not verify a mechanism underlying the present survey geometry.

Interventions also have limitations. Cultural

steering dimensions can be entangled (Dang and Masud, 2026); sparse-autoencoder steering does not consistently outperform prompting or fine-tuning in benchmark comparisons (Wu et al., 2025). Feature absorption (Chanin et al., 2025) and predictable reconstruction error (Engels et al., 2025) complicate feature-based explanations. The field’s stocktake accordingly describes mechanistic interpretability as promising rather than established behavioural assurance (Sharkey et al., 2025). Such methods may test hypotheses suggested by these results, not certify cultural neutrality.

Stability depends on the task and statistic.

Moore et al. (2024) find relative consistency across paraphrase, format and translation, with base models more consistent than fine-tuned ones in their comparisons. Kovač et al. (2024) distinguish model-dependent ipsative stability without persona instructions from low rank-order stability of simulated personas; the human benchmarks are not directly equivalent. Huang (2026) separates order-and-wording effects from an inferred underlying stance in their moral-judgement task. These studies motivate separating response-style variation from position estimates, but our constructed-profile variance summaries do not identify the same latent stance or establish comparable stability.