

Unmasking the Factual-Conceptual Gap in Persian Language Models

Alireza Sakhaeirad*
EPFL

Ali Ma'manpoosh
University of Isfahan

Arshia Hemmat
University of Oxford

Abstract

While emerging Persian NLP benchmarks have expanded into pragmatics and politeness, they rarely distinguish between memorized cultural facts and the ability to reason about implicit social norms. We introduce DIVANBENCH, a diagnostic benchmark focused on superstitions and customs, arbitrary, context-dependent rules that resist simple logical deduction. Through 315 questions across three task types (factual retrieval, paired scenario verification, and situational reasoning), we evaluate seven Persian LLMs and reveal three critical failures: most models exhibit severe acquiescence bias, correctly identifying appropriate behaviors but failing to reject clear violations; continuous Persian pretraining amplifies this bias rather than improving reasoning, often degrading the model's ability to discern contradictions; and all models show a 21% performance gap between retrieving factual knowledge and applying it in scenarios. These findings demonstrate that cultural competence requires more than scaling monolingual data, as current models learn to mimic cultural patterns without internalizing the underlying schemas.¹

1 Introduction

If you offer a Persian guest food, they will refuse. If you offer again, they will refuse. Only on the third offer, the “real” one, will they accept. This three-iteration ritual, called *taarof*, is immediately obvious to any Iranian child raised in the culture. But can a language model, trained on billions of Persian tokens, distinguish genuine *taarof* from a cultural violation? Our results suggest: not really.

The rapid advancement of large language models (LLMs) has sparked significant interest in their multilingual capabilities, including Persian language

*Correspondence to: alireza.sakhaeirad@epfl.ch

¹Dataset publicly available huggingface.co/datasets/divanbench/divanbench

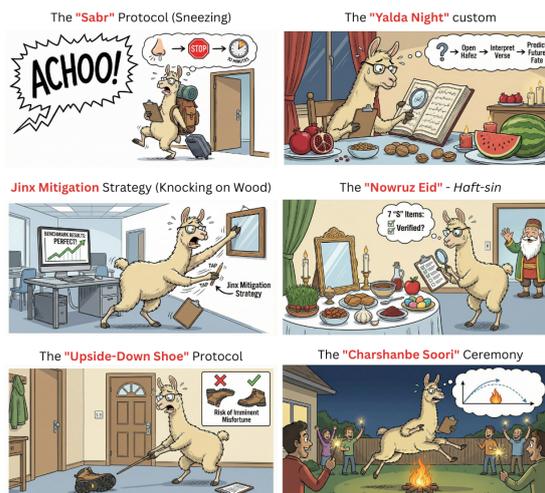


Figure 1: Sample Persian cultural concepts from the benchmark, spanning superstitions, traditions, and taboos.

processing. Current Persian LLM benchmarks primarily focus on factual knowledge retrieval, translation quality, and linguistic tasks such as sentiment analysis and named entity recognition (Khashabi et al., 2020). However, these evaluations critically overlook a dimension central to Persian communication: the ability to reason about implicit cultural concepts, particularly superstitions and customs, that govern appropriate behavior in context-dependent social situations.

Persian culture embeds meaning across multiple historical layers: ancient Zoroastrian beliefs (*Chaharshanbe Suri* fire-jumping, *esfand* burning for evil eye protection), Islamic traditions (*nazri* vow offerings, *ghorbani* sacrifice), complex social etiquette systems (*taarof*, *jang-e hesab*), rich folk cosmology featuring supernatural beings (*jinn*, *div*, *pari*, *bakhtak*), widespread superstitious practices (whistling at night attracts evil), and elaborate life-cycle ceremonies with highly specific ritualistic requirements. These concepts are not facts to be retrieved from a knowledge base; they are cultural

Category	Example Concepts
Social Etiquette	<i>Taarof, Jang-e Hesab, Doorway Deference, Three-Times Rule, Pishkesh, Shirini</i>
Nowruz Traditions	<i>Haft Sin, Chaharshanbe Suri, Haji Firuz, Samanu Pazan, Khaneh Tekani, Eidi</i>
Supernatural Beings	<i>Jinn, Div, Pari, Bakhtak, Hamzad, Al</i>
Apotropaic Rituals	<i>Nazar amulet, Esfand burning, Salt circle, Bismillah invocations</i>
Wedding Ceremonies	<i>Sofreh Aghd, Knife Dance, Kuzeh Shekani, Honey ritual</i>
Taboos	Whistling at night, Sweeping at night, Stepping on bread, Shoe taboos, Nail cutting restrictions
Divination/Omens	<i>Fal-e Hafez, Fal-Gush, Dream interpretation, Itchy palms, Ear ringing</i>

Table 1: Distribution of cultural concepts across categories. Concepts span ancient Zoroastrian beliefs, Islamic traditions, social etiquette systems, folk cosmology, life-cycle ceremonies, taboos, and divination practices.

schemas (Shore, 1996; Strauss and Quinn, 1997) requiring understanding of context-dependent social dynamics, implicit power relations, and culturally-specific rules of appropriateness that are obvious to cultural insiders but opaque to pattern-matchers.

1.1 Motivation: Why Superstitions and Customs?

Superstitions and customs provide particularly demanding testbeds for cultural reasoning because they often lack logical justification-, unlike social etiquette, which has pragmatic benefits, superstitions require acceptance of culturally-transmitted beliefs without empirical grounding. Furthermore, they exhibit context-dependence: the same action (such as sweeping) is neutral by day but taboo at night. These practices involve implicit rules that are typically corrected through social feedback rather than explicit instruction, and Persian superstitions notably reflect historical syncretism, layering Zoroastrian, Islamic, and folk beliefs in complex ways. Recent work on Korean superstitions (Kim and Lee, 2025) demonstrated that culture-bound phenomena reveal limitations invisible in standard benchmarks. We extend this insight to Persian, where superstitions permeate daily life and provide rich signal for distinguishing genuine cultural understanding from keyword-matching.

1.2 Research Questions, Contributions, and Key Findings

This work (illustrated in Fig. 1) investigates whether Persian-capable Large Language Models (LLMs) internalize *cultural schemas*, the implicit social logic mapping context to action, or merely reproduce surface-level *cultural facts* and stereotypes. The investigation focuses on three core dimensions: the propensity for models to exhibit acquiescence bias through pattern-matching, the impact of continuous Persian pretraining on reasoning versus fluency, and the capacity for factual knowledge to be operationalized into scenario-based application. By

distinguishing between fluent cultural expression and robust cultural reasoning, this work provides a framework for evaluating cultural competence in low-resource linguistic environments.

Central to this evaluation is the introduction of DIVANBENCH, a diagnostic framework covering 81 concepts across superstitions, customs, and social etiquette, instantiated through 315 questions. The benchmark adopts a three-level architecture designed to isolate distinct layers of cultural competence: Factual MCQ for baseline knowledge retrieval, Binary Belief Verification for testing discernment through paired positive/negative scenarios, and Scenario-Based MCQ for evaluating multi-step inference in complex social contexts. By utilizing a paired design, contrasting behavior that aligns with Persian customs against plausible but culturally inappropriate alternatives, this methodology effectively quantifies acquiescence bias. Furthermore, the study isolates the effects of monolingual adaptation by comparing a base model, Llama 3.1-8B, its Persian-adapted variant Dorna2-8B, revealing how specialized pretraining affects the transition from surface fluency to logical consistency.

Our evaluation of 7 state-of-the-art models yields the following key findings:

- **The acquiescence Trap:** 6 of 7 models exhibit significant asymmetry in reasoning; while they identify appropriate behaviors with 84%–92% accuracy, they fail to reject cultural violations at rates between 19% and 48%, suggesting a reliance on pattern-matching over logic.
- **The Persian Pretraining Paradox:** Continuous pretraining can inadvertently degrade critical reasoning. Llama 3.1-8B’s rejection accuracy of 73% dropped to 30% after Persian-specific adaptation (Dorna2), representing a 43-percentage-point decline.
- **The Factual–Conceptual Gap:** Performance

decreases by an average of 21% when transitioning from factual retrieval to scenario-based reasoning, demonstrating that memorized cultural facts do not reliably translate into functional cultural schemas.

- **Distinct Learning Targets:** The results provide empirical evidence that cultural facts and schemas behave as independent objectives, suggesting that cultural competence requires training mechanisms beyond standard data scaling or monolingual adaptation.

The remainder of this paper is structured as follows: Section 2 reviews related work on cultural evaluation and Persian NLP; Section 3 describes our benchmark design, cultural concept coverage, and data collection methodology. Section 4 details our experimental setup including model selection, evaluation protocol, and experimental design. Subsection 4.2 presents our four main findings on acquiescence bias, pretraining effects, factual-conceptual gaps, and model scaling. Subsection 4.3 analyzes the mechanisms behind these findings and their implications for multilingual NLP. Finally, Section 5 concludes and discusses future directions for cultural reasoning research. Appendix A gives more insight into the data and presents more examples, and appendix B adds some additional detail on system prompt variations, experiments, and results.

2 Related Work

2.1 Persian LLM Evaluation

Most Persian NLP benchmarks have traditionally emphasized linguistic competence and factual or domain knowledge. Representative resources include PARSINLU for broad NLU coverage (Khashabi et al., 2020), FARSTAIL for Persian natural language inference (Amirkhani et al., 2021), and PERSIANMEDQA for bilingual medical question answering (Ranjbar Kalahroodi et al., 2025). Recent efforts broaden evaluation toward general knowledge and multimodal educational assessment, including the Khayyam Challenge (PersianMMLU) (Ghahroodi et al., 2024) and MEENA (PersianMMMU) (Ghahroodi et al., 2025). In parallel, a newer line of work targets cultural and pragmatic competence more directly: PERCUL uses story-driven scenarios to probe cultural understanding in Persian (Moosavi Monazzah et al., 2025), TAAROFBENCH evaluates the Persian politeness system and its nuanced social dynam-

ics (Gohari Sadr et al., 2025), and ELAB benchmarks Persian-relevant alignment and safety dimensions (Pourbahman et al., 2025). Together, these works show a shift from “language + facts” evaluation toward culturally grounded pragmatics, though the latter remains less standardized and comparatively underexplored.

2.2 Cultural Competence in LLMs

Beyond Persian, substantial evidence shows that LLMs often fail to capture culture-specific, context-dependent norms and may default to dominant viewpoints or produce plausible-but-shallow cultural explanations (Zhang et al., 2025; Dai et al., 2025; Durmus et al., 2024; Hossain and Afli, 2025). Cross-cultural evaluations similarly report performance disparities between high-resource and low-resource cultures and uneven representation of global perspectives (Cao et al., 2023; Durmus et al., 2023). Related research in social/pragmatic reasoning and perspective-taking further indicates weaknesses on tasks that require implicit modeling of beliefs and intentions (Sap et al., 2019, 2022). Particularly relevant are cultural testbeds based on superstitions (e.g., NUNCHI-BENCH), where models may handle surface facts but struggle with situational reasoning (Kim and Lee, 2025).

Methodologically, evaluation can be confounded by systematic response biases: LLMs exhibit position bias in multiple-choice settings and related tendencies such as acquiescence bias/“yes-saying” (Zheng et al., 2023; Wang et al., 2023; Wallace et al., 2019; Si et al., 2023). Prior work mitigates these effects via prompt design and calibration (Ko et al., 2023; Si et al., 2023), while alternative evaluation designs aim to measure bias explicitly. Finally, cognitive theories distinguish explicit cultural facts from implicit cultural schemas learned through repeated, socially situated experience (Shore, 1996; Strauss and Quinn, 1997). This distinction aligns with arguments that text-only training may be insufficient for grounded understanding (Bisk et al., 2020) and echoes the classic observations that robust common sense reasoning is difficult to recover from text statistics alone (Levesque et al., 2012).

2.3 Positioning Our Work

Our work uniquely combines: (1) comprehensive coverage of Persian superstitions and customs (more than 80 concepts across 7 domains); (2) explicit bias measurement through paired pos-

itive/negative design; (3) direct comparison isolating pretraining effects (Llama3.1 (Dubey et al., 2024) vs. Dorna2 (PartAI Team, 2024)); and (4) theoretical grounding distinguishing cultural facts from schemas. While prior work addresses individual aspects (Persian-related factual retrieval tests, cultural evaluation, low-resource pretraining), we provide integrated analysis revealing systematic limitations in current approaches to cultural competence in Persian LLMs.

3 Benchmark Design and Data Collection

We introduce DIVANBENCH, a benchmark for evaluating cultural reasoning in Persian LLMs. Our design philosophy centers on distinguishing genuine cultural understanding from pattern-matching to plausible-sounding Persian text.

3.1 Design Principles

We evaluate models across three complementary tasks isolating different aspects of cultural competence:

Factual Multiple-Choice Questions The first data type consists of 100 Factual MCQs designed to test whether models can accurately retrieve fundamental information regarding Persian culture, history, geography, and traditions. These questions serve as a critical baseline control, identifying whether a model possesses the raw cultural knowledge necessary for more complex tasks. Success in this category indicates a robust factual foundation, while failure suggests a lack of exposure to Persian-specific data during the model’s pre-training phase.

Binary Belief Verification The second task type comprises 162 statements categorized into matched pairs for 81 distinct cultural concepts. Each pair includes a "positive" scenario indicating a person following the Persian customs and a "negative" scenario where the subject acts against those customs in a way that remains plausible in a non-Persian context. This paired design is specifically intended to measure acquiescence bias and discernment; models that rely on simple pattern matching often show an asymmetric performance, erroneously accepting culturally inappropriate behaviors while correctly identifying appropriate ones.

Scenario-Based Multiple-Choice Questions

The final data type includes 53 Scenario-Based MCQs that challenge models to apply cultural reasoning to complex, novel social situations.

These scenarios require intricate inferences about social hierarchies, contextual appropriateness, and interpersonal nuances that are not explicitly taught. By focusing on concepts typically acquired through lived experience and social interaction, this task isolates the model’s ability to move beyond rote memorization and demonstrate authentic cultural competence in dynamic settings.

Example: Scenario-Based MCQ

Arman has a sore throat and a high fever. His mother brings him soup and tea but strictly warns him not to eat the fresh melon sitting on the counter. Arman argues that the fruit is full of vitamins and will help him recover. Why does the mother forbid Arman from eating the melon while he has a cold?

- (A) Melon is *Sard* (Cold); adding cold fruit while sick will freeze the lungs and prolong recovery
- (B) Melon is *Garm* (Hot); eating it during a fever will ignite the blood, spreading infection faster
- (C) Folk belief holds melon aroma attracts night-illness spirits, preventing medicine from working
- (D) Hygiene precaution: high sugar content feeds bacteria in the throat, causing permanent voice loss

Expected: (A) — requires understanding the classical Persian food classification system (*Sard-Garm*) and traditional Persian medicine, which lacks modern medical evidence.

Author-Generated Content and Quality Control

All questions were drafted and refined through multiple rounds of author review based on cultural knowledge acquired via lived experience in Iranian society. This process ensures authenticity by reflecting real social practices rather than stereotypes, focusing on "insider knowledge" that is intuitive to cultural members but requires active reasoning from outsiders. We prioritized implicit concepts typically learned through social interaction, rather than explicit instruction. By ensuring unambiguity for insiders while maintaining plausible alternatives for outsiders, the dataset enables a diagnostic evaluation of cultural competence that distinguishes genuine understanding from simple pattern

matching.

3.2 Cultural Concept Coverage

Our benchmark spans seven major cultural domains, summarized in Table 1 and detailed in Appendix A.1. These domains encompass superstitions and omens such as divination practices (*Fal-e Hafez*, *Fal-Gush*), dream interpretation, and bodily omens like itchy palms or ear ringing, alongside apotropaic rituals like burning *esfand* to ward off the evil eye, wearing *nazar* amulets, and *Bismillah* invocations. Additionally, the dataset covers taboos governing daily life, including restrictions on whistling or sweeping at night, and prohibitions against stepping on bread.

4 Experiments

4.1 Setup

Model Selection We evaluate 7 models from the Open Persian LLM Leaderboard (OpenPersian Team, 2024) with similar parameter counts to ensure fair comparison. All models are in the 7–12B parameter range, representing state-of-the-art performance for Persian language tasks: Aya-8B (Co-here For AI, 2024), Dorna2-8B (PartAI Team, 2024), Gemma2-9B and Gemma3-12B (Gemini Team, 2023), Llama3.1-8B (Dubey et al., 2024), and Qwen2-7B and Qwen2.5-7B (Qwen Team, 2024). The critical comparison is Llama3.1-8B (base) versus Dorna2-8B (Persian-adapted through continuous pretraining), which isolates the effect of Persian-focused pretraining while controlling for all other modeling choices. (For detailed description about models see App. B.1)

Evaluation Protocol Following best practices for reproducible LLM evaluation (Chang et al., 2024), we use temperature of 0.1 and top-p sampling set to 0.9 with fixed random seeds. For answer extraction, we employ GPT-4o-mini as a systematic extraction agent (Zheng et al., 2023), which parses model outputs to identify selected options (A/B/C/D) for multiple-choice questions and yes/no responses for binary questions. To account for prompt sensitivity (Si et al., 2023), each question is evaluated with 5 diverse system prompts that vary in phrasing while maintaining semantic equivalence. All prompts instruct models to respond as Iranian cultural insiders to ensure fair evaluation conditions. We report mean accuracy and standard deviation across prompt variations. To see exact system prompts, see B.2.

We compute three complementary metrics: (1) **Accuracy** measuring percentage of correct responses per task type; (2) **Acquiescence bias** quantifying the difference between acceptance rates for positive scenarios and rejection rates for negative scenarios in binary tasks, where high positive bias indicates pattern-matching rather than reasoning; and (3) **Factual-Conceptual Gap** measuring performance difference between factual retrieval and scenario-based reasoning. These metrics provide multi-dimensional assessment of cultural competence beyond standard accuracy reporting.

Experimental Design We conduct three complementary experiments to systematically evaluate cultural reasoning in Persian LLMs:

- 1. Acquiescence bias Measurement (Experiment 1).** We compare model performance on paired positive and negative binary scenarios testing identical cultural concepts. This paired design directly measures acquiescence bias as $B = \text{Acc}_{\text{pos}} - \text{Acc}_{\text{neg}}$, where Acc_{pos} is accuracy on appropriate behavior and Acc_{neg} is accuracy on violations. Models with high positive bias demonstrate pattern-matching to cultural keywords rather than genuine reasoning.
- 2. Pretraining Effects (Experiment 2).** We conduct a controlled comparison between Llama3.1-8B (base model M_{base}) and Dorna2-8B (Persian-adapted model M_{adapted}). This matched-pair design isolates the effect of Persian pretraining by comparing performance changes $\Delta = \text{Acc}(M_{\text{adapted}}) - \text{Acc}(M_{\text{base}})$ across all task types, revealing whether continuous pretraining improves cultural reasoning or reinforces surface-level pattern-matching.
- 3. Knowledge Transfer (Experiment 3).** We analyze the factual-conceptual gap as $G = \text{Acc}_{\text{factual}} - \text{Acc}_{\text{scenario}}$, where $\text{Acc}_{\text{factual}}$ measures retrieval and $\text{Acc}_{\text{scenario}}$ measures cultural reasoning. Large consistent gaps suggest cultural facts and schemas engage different cognitive mechanisms.

4.2 Results

We present results across all evaluation dimensions, revealing systematic patterns in how Persian language models handle cultural reasoning. Table 2

Model	Factual MCQ (Knowledge)	Scenario MCQ (Reasoning)	Gap (Fact - Scenario)	Acquiescence bias (Accept - Reject)
Gemma3-12B	87.0 ± 1.0	66.0 ± 1.0	21.0	+61.0
Gemma2-9B	83.0 ± 0.0	65.0 ± 3.0	18.0	+41.0
Llama3.1-8B	80.0 ± 2.0	53.0 ± 5.0	27.0	-10.0
Qwen2.5-7B	75.0 ± 1.0	60.0 ± 2.0	15.0	+26.0
Qwen2-7B	73.0 ± 3.0	48.0 ± 5.0	25.0	+65.0
Dorna2-8B	73.0 ± 2.0	54.0 ± 3.0	19.0	+61.0
Aya-8B	66.0 ± 2.0	42.0 ± 2.0	24.0	+19.0
Average	76.7	55.4	21.3	+43.3

Table 2: Core performance metrics across all models. **Gap** measures difficulty transferring factual knowledge to scenario reasoning. **Acquiescence bias** (Accept - Reject) reveals pattern-matching: positive values (red) indicate models accept appropriate behavior readily but fail to reject violations. Only Llama3.1-8B shows skeptical bias (blue). Percentages shown; standard deviations computed across 5 prompt variations. Models sorted by Factual MCQ performance.

Task	Llama3.1-8B	Dorna2-8B	Change
Reject False	0.73	0.30	-43%
Accept True	0.63	0.91	+28%
Factual MCQ	0.80	0.73	-7%
Scenario MCQ	0.53	0.54	+1%

Table 3: Direct comparison of Llama3.1-8B (base) and Dorna2-8B (Persian-tuned). Persian pretraining dramatically increased acquiescence while destroying critical reasoning.

shows complete performance across all tasks (See table 5 for detailed results of binary tests).

Finding 1: The acquiescence Trap Most models exhibit severe acquiescence bias on cultural questions, accepting plausible statements far more readily than rejecting violations. This asymmetry reveals that models are pattern-matching to culturally-themed Persian text rather than reasoning about appropriateness. Extreme cases include Qwen2-7B (84% acceptance vs. 19% rejection, bias = +65%), Gemma3-12B (92% acceptance vs. 31% rejection, bias = +61%), and Dorna2-8B (91% acceptance vs. 30% rejection, bias = +61%). These models correctly identify culturally appropriate behavior when presented positively but fail to recognize violations of the same concepts when presented negatively. For instance, accepting both "Standing up when elders arrive" (correct code of respect) and "Remaining seated for comfort when elders arrive" (violates code of respect). The only exception is Llama3.1-8B, showing opposite bias (73% rejection vs. 63% acceptance, bias = -10%), suggesting an skeptical tone in this model.

Finding 2: The Persian Pretraining Paradox Comparing Llama3.1-8B (base) with Dorna2-8B (Persian-tuned) in table 3 reveals counterintuitive



Figure 2: acquiescence bias across models. Most models accept true cultural statements readily (Blue) but fail to reject false ones (Red), indicating pattern-matching rather than reasoning.

effects of continuous pretraining on cultural reasoning. Rejection accuracy collapsed dramatically from 73% to 30% (-43 percentage points), while acceptance accuracy soared from 63% to 91% (+28pp). Notably, factual knowledge slightly declined (80% to 73%) and scenario reasoning remained essentially flat (53% to 54%), indicating that continuous Persian pretraining taught the model to recognize cultural patterns without improving reasoning about cultural logic. This suggests that more Persian data reinforced surface-level pattern-matching over deep cultural understanding. Dorna2 became more "culturally compliant" but less discerning, accepting any plausible-sounding cultural scenario regardless of correctness. This finding challenges the common assumption in low-resource NLP that monolingual pretraining universally improves understanding.

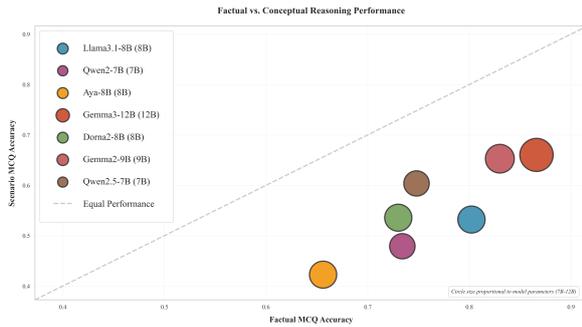


Figure 3: Factual Knowledge vs. Cultural Reasoning Gap. Gemma3-12B (largest bubble, top-right) achieves highest factual accuracy but shows inconsistent scenario reasoning. All models fall below the diagonal, indicating systematic difficulty in transferring factual knowledge to cultural schema application. Bubble size represents model parameters.

Finding 3: The Factual-Conceptual Gap All models show consistent performance degradation from factual to scenario-based tasks, with an average gap of 21 percentage points. Models can retrieve facts about Persian culture ("What is *taarof*?") but struggle to apply cultural logic to novel scenarios ("In this situation, would *taarof* be appropriate?"). This gap demonstrates that knowing that a concept exists differs fundamentally from knowing when it appropriately applies. The gap persists even for the best-performing model (Gemma3-12B shows -21%), suggesting this is a fundamental limitation rather than a simple capacity issue. This evidence supports the cognitive anthropology distinction between cultural facts (retrievable from text) and cultural schemas (requiring embodied social learning).

Finding 4: Size Does Not Guarantee Cultural Intelligence Gemma3-12B, the only model exceeding 10B parameters, shows contradictory performance patterns: it achieves best factual retrieval (87%, 21% above average) yet worst acquiescence bias (+61%, tied with Dorna2). The larger model excels at memorizing cultural facts while showing extreme acquiescence bias, accepting 92% of positive statements but rejecting only 31% of negative ones. This suggests that model size improves fact memorization without improving cultural discernment, and larger models may amplify pattern-matching behaviors learned from training data.

4.3 Analysis

Why Persian Cultural Concepts Are Hard Persian cultural concepts exhibit three critical char-

acteristics that resist pattern-matching. First, they encode implicit social contracts never stated in text: *taarof* requires tracking iteration counts, modeling power asymmetries, and distinguishing ritual from sincerity through multi-turn interaction. Second, they involve context-dependent inversion where identical actions have opposite meanings: paying for a meal is appropriate among peers after negotiation but offensive when done secretly with elders. Third, they conflate factual and normative knowledge: "Is burning *esfand* effective?" (factual, no) versus "Is burning *esfand* appropriate when moving?" (normative, yes). Models trained on factual text systematically conflate these categories.

The Acquiescence Trap Mechanism The severe acquiescence bias across six models stems from distributional patterns in Persian training data. Models learn that cultural keywords like *taarof*, *nazri*, and *esfand* co-occur with positive contexts: respectful news coverage, affirmative social media posts, and celebratory descriptions. When encountering these keywords, models trigger acceptance regardless of scenario details. This explains why they accept both correct *taarof* (refuse twice, accept third time) and violations (accept immediately). Without reasoning about iteration logic, they pattern-match to cultural keywords instead.

Why Persian Pretraining Degraded Reasoning Continuous Persian pretraining on Dorna2 collapsed rejection accuracy by 43 percentage points while boosting acceptance by 28 points. This asymmetry reveals surface pattern reinforcement: Persian corpora contain overwhelmingly affirmative mentions of traditions, creating distributional bias toward accepting culturally-themed statements. Meanwhile, Llama3.1's skeptical stance (73 percent rejection) likely reflects instruction tuning for critical evaluation. Persian pretraining overrode this skepticism, producing cultural fluency without reasoning. This challenges the assumption that monolingual pretraining improves understanding, it may amplify fluency while degrading discernment.

Factual-Conceptual Gap as Schema Evidence The consistent 21 percent performance drop from factual to scenario tasks operationalizes the distinction between cultural facts (retrievable knowledge like "Nowruz marks New Year") and cultural schemas (situational rules like "Would serving *halva* be appropriate here?"). Facts require memorization; schemas require conditional reasoning

about context, relationships, and implicit norms. Current training successfully instills facts but fails at schemas, suggesting that adding more factual data cannot close this gap, different learning mechanisms are needed.

5 Conclusion

We introduced DIVANBENCH, a benchmark evaluating cultural reasoning in Persian LLMs through 81 concepts and 315 questions across three task types. Evaluation of seven models reveals four key findings: most models exhibit severe acquiescence bias, accepting appropriate behavior while failing to reject violations; continuous Persian pre-training amplified this bias rather than improving reasoning; a consistent 21 percent gap between factual and scenario performance demonstrates that cultural facts and schemas engage different mechanisms; and model size improves memorization but not reasoning. These findings challenge assumptions that scaling monolingual data improves cultural competence and suggest that low-resource languages require more advanced methods, rather than simply training the model on more tokens of the target language. We release our benchmark publicly to enable systematic measurement of cultural understanding, providing a template for evaluation in other languages. Persian culture’s synthesis of Zoroastrian, Islamic, and folk traditions makes it a demanding testbed for genuine schema reasoning versus surface pattern-matching. Our work establishes a foundation for building language models that truly comprehend the cultures they serve.

Limitations

Model size and scaling regime. Our study evaluates a narrow band of *small-to-mid sized* open models (7–12B parameters), selected to enable controlled comparisons across systems with broadly similar capacity and inference cost. This design improves comparability, but it limits what we can conclude about how cultural conceptual reasoning behaves at larger scales (e.g., 30B–70B+), where instruction tuning, longer context windows, or different training mixtures may change both *bias profiles* and *reasoning strategies*. In particular, our observation that models can improve factual recall while still exhibiting accept-over-reject asymmetry may not extrapolate monotonically to substantially larger model families; scaling could either attenuate or amplify the “acquiescence trap,” depending on how the model is trained and aligned. Additionally, by keeping model size relatively constant, we cannot disentangle whether certain failures are fundamentally capacity-limited versus primarily data- and objective-driven.

Hand-curated dataset and author priors. DivanBench is authored manually based on lived cultural knowledge, which increases scenario realism and insider validity but introduces several curation constraints. First, manual writing inevitably reflects the authors’ judgments about what counts as the “canonical” interpretation of a concept, which may underrepresent regional, socioeconomic, or generational variation. Second, even with iterative review, scenario phrasing can unintentionally leak cues (e.g., overly salient keywords, unnatural dialogue) that models may exploit via superficial heuristics; especially in paired positive/negative formats. Third, manual coverage is bounded: although we span many (more than 80) concepts, the space of culturally meaningful edge cases is far larger, and some concepts may be overrepresented by the scenarios that are easiest to write unambiguously. Finally, because expert annotation is expensive and time-consuming, we do not incorporate large-scale crowdsourced validation or inter-annotator agreement; thus, while items are intended to be unambiguous to cultural insiders, residual ambiguity cannot be ruled out.

References

- Hossein Amirkhani, Mohammad AzariJafari, Zohreh Pourjafari, Soroush Faridan-Jahromi, Zeinab Kouhkan, and Azadeh Amirak. 2021. [Farstail: A persian natural language inference dataset](#). *Preprint*, arXiv:2009.08820.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, and 1 others. 2020. Experience grounds language. *Proceedings of EMNLP*.
- Yong Cao, Li Zhou, Seolhwa Lee, Laura Cabello, Min Chen, and Daniel Hershcovich. 2023. Assessing cross-cultural alignment between chatgpt and human societies: An empirical study. *arXiv preprint arXiv:2303.17466*.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, and 1 others. 2024. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*.
- Cohere For AI. 2024. Aya model: An instruction finetuned open-access multilingual language model. *arXiv preprint arXiv:2402.07827*.
- 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](#). *Preprint*, arXiv:2505.18562.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Esin Durmus, Karina Nguyen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, and 1 others. 2024. [Towards measuring the representation of subjective global opinions in language models](#). *Preprint*, arXiv:2306.16388.
- Esin Durmus, Karina Nguyen, Thomas I. Liao, Nicholas Schiefer, and 1 others. 2023. [Towards measuring the representation of subjective global opinions in language models](#). *arXiv preprint arXiv:2306.16388*.
- Gemini Team. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Omid Ghahroodi, Arshia Hemmat, Marzia Nouri, Seyed Mohammad Hadi Hosseini, Doratossadat Dastgheib, Mohammad Vali Sanian, Alireza Sahebi, Reihaneh Zohrabi, Mohammad Hossein Rohban, Ehsaneddin Asgari, and Mahdieh Soleymani Baghshah. 2025. [Meena \(persianmmmu\): Multimodal-multilingual educational exams for n-level assessment](#). *Preprint*, arXiv:2508.17290.
- Omid Ghahroodi, Marzia Nouri, Mohammad Vali Sanian, Alireza Sahebi, Doratossadat Dastgheib, Ehsaneddin Asgari, Mahdieh Soleymani Baghshah, and Mohammad Hossein Rohban. 2024. [Khayyam challenge \(persianmmlu\): Is your llm truly wise to the persian language?](#) *Preprint*, arXiv:2404.06644.
- Nikta Gohari Sadr, Sahar Heidariasl, Karine Megerdoo-mian, Laleh Seyyed-Kalantari, and Ali Emami. 2025. [We politely insist: Your llm must learn the persian art of taarof](#). *Preprint*, arXiv:2509.01035.
- Shehenaz Hossain and Haithem Afi. 2025. [Craft: An explanation-based framework for evaluating cultural reasoning in multilingual language models](#). *Preprint*, arXiv:2510.14014.
- Daniel Khashabi, Arman Cohan, Siamak Shakeri, Pedram Hosseini, Pouya Pezeshkpour, Graham Neubig, and 1 others. 2020. Parsinlu: A suite of language understanding challenges for persian. *arXiv preprint arXiv:2012.06154*.
- Kyuhee Kim and Sangah Lee. 2025. [Nunchi-bench: Benchmarking language models on cultural reasoning with a focus on Korean superstition](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 15328–15342, Vienna, Austria. Association for Computational Linguistics.
- Yoonjoo Ko, Hyungjoo Lee, Taehwan Lee, Jaewook Jung, Jinhyeok Park, and Jaihoon Choi. 2023. [From language modeling to instruction following: Understanding the behavior shift in llms after instruction tuning](#). *arXiv preprint arXiv:2310.00492*.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Proceedings of KR*.
- Erfan Moosavi Monazzah, Vahid Rahimzadeh, Yadollah Yaghoobzadeh, Azadeh Shakery, and Mohammad Taher Pilehvar. 2025. [Percul: A story-driven cultural evaluation of llms in persian](#). *Preprint*, arXiv:2502.07459.
- OpenPersian Team. 2024. [Openpersian leaderboard: Benchmarking persian language models](#). <https://huggingface.co/spaces/OpenPersian/leaderboard>.
- PartAI Team. 2024. [Dorna: Persian-optimized language models](#). <https://huggingface.co/PartAI/Dorna-Llama3-8B-Instruct>.
- Zahra Pourbahman, Fatemeh Rajabi, Mohammadhossein Sadeghi, Omid Ghahroodi, Somaye Bakhshaei, Arash Amini, Reza Kazemi, and Mahdieh Soleymani Baghshah. 2025. [Elab: Extensive llm alignment benchmark in persian language](#). *Preprint*, arXiv:2504.12553.
- Qwen Team. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Mohammad Javad Ranjbar Kalahroodi, Amirhossein Sheikholeslami, Sepehr Karimi, Sepideh Ranjbar Kalahroodi, Hessaam Faili, and Azadeh Shakery. 2025. [Persianmedqa: Evaluating large language models on a persian-english bilingual medical question answering benchmark](#). *Preprint*, arXiv:2506.00250.

- Maarten Sap, Ronan LeBras, Daniel Fried, and Yejin Choi. 2022. Neural theory-of-mind? on the limits of social intelligence in large llms. *arXiv preprint arXiv:2210.13312*.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. 2019. Social iqa: Commonsense reasoning about social interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, pages 4463–4473.
- Bradd Shore. 1996. *Culture in Mind: Cognition, Culture, and the Problem of Meaning*. Oxford University Press.
- Chenglei Si, Zihan Gan, Zhengyuan Yang, Shuohang Wang, Jingfeng Wang, Jordan Boyd-Graber, and Lijuan Wang. 2023. Prompting gpt-3 to be reliable. In *Proceedings of ICLR*.
- Claudia Strauss and Naomi Quinn. 1997. *A Cognitive Theory of Cultural Meaning*. Cambridge University Press.
- Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan Boyd-Graber. 2019. Trick me if you can: Human-in-the-loop generation of adversarial examples for question answering. In *Transactions of the Association for Computational Linguistics*, volume 7, pages 387–401.
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Zhifang Sui, and 1 others. 2023. Label words are anchors: An information flow perspective for understanding in-context learning. In *Proceedings of EMNLP*, pages 9840–9855.
- Jinghao Zhang, Sihang Jiang, Shiwei Guo, Shisong Chen, Yanghua Xiao, Hongwei Feng, Jiaqing Liang, Mingguo HE, Shimin Tao, and Hongxia Ma. 2025. *Culturescope: A dimensional lens for probing cultural understanding in llms*. *Preprint*, arXiv:2509.16188.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*.

A Dataset Details

A.1 Detailed Cultural Concept Lists

Table 4 provides comprehensive concept coverage across all seven cultural domains.

A.2 Additional Examples

Example: Factual MCQ

What fruit is used as a symbol on the Yalda night table?

(A) Orange (B) Apple (C) Pomegranate (D) Pear

Expected: (C) Pomegranate

Example: Binary Belief Verification

Positive: “During a casual family gathering, an elderly aunt walks into the room. Everyone younger immediately stands up as a respectful acknowledgment of her presence. Did they act according to Persian tradition?”

Expected: Yes (This follows the rule to stand up when elders enter)

Negative: “During a family gathering, an elderly aunt enters the room. Since everyone is already comfortable, no one stands up and they simply greet her warmly from their seats. Did they act according to Persian tradition?”

Expected: No (This violates the rule to stand up when elders enter)

B Appendix

B.1 Model Descriptions

Aya-8B Multilingual model from CohereForAI with Persian support. Trained on diverse multilingual data including Persian web text and translated datasets. Instruction-tuned for helpfulness.

Dorna2-8B Persian-optimized variant of Llama3.1-8B created through continuous pretraining on Persian corpora including news, social media, literature, and web text. Maintained Llama3.1’s instruction-tuning while adapting vocabulary and linguistic patterns to Persian.

Gemma2-9B and Gemma3-12B Google’s open-source multilingual foundation models (second and third generations). Trained on diverse web data including Persian content. Gemma3 represents scaled-up version with 12B parameters.

Category	Example Concepts
Social Etiquette	<i>Taarof</i> (ritual politeness), <i>Jang-e Hesab</i> (payment battles), Doorway Deference (older enters first), Three-Times Rule (refuse twice), <i>Pishkesh</i> (gift-giving protocol), <i>Shirini</i> (sweet bringing obligation), Guest-host dynamics, Shoe removal customs, Greeting hierarchies
Superstitions & Omens	<i>Fal-e Hafez</i> (Hafez divination), <i>Fal-Gush</i> (eavesdropping for omens), Dream interpretation, Itchy palms (money coming), Ear ringing (someone talking about you), Twitching eye (bad omen), Sneezing omens, Bird flight patterns, Broken mirrors, Spilled salt
Apotropaic Rituals	<i>Esfand</i> burning (wild rue fumigation), <i>Nazar</i> amulet (evil eye protection), Salt circle protection, <i>Bismillah</i> invocations, Garlic hanging, Iron deterrence, Seven knots ritual, Mirror placement, Knife under pillow, Holy verses
Supernatural Beings	<i>Jinn</i> (fire spirits), <i>Div</i> (demons), <i>Pari</i> (fairies), <i>Bakhtak</i> (sleep paralysis demon), <i>Hamzad</i> (personal spirit double), <i>Al</i> (child-stealing demon), <i>Ghoul</i> , Evil eye personification, Ancestor spirits
Nowruz Traditions	<i>Haft Sin</i> (seven S’s table), <i>Chaharshanbe Suri</i> (fire-jumping), <i>Haji Firuz</i> (blackface herald), <i>Samanu Pazan</i> (wheat pudding stirring), <i>Khaneh Tekani</i> (spring cleaning), <i>Eidi</i> (new year gift), <i>Sizdah Bedar</i> (13th day picnic), <i>Sabzeh</i> growing, <i>Goldfish</i> symbolism, New clothes tradition, Elder visitation order, Fire-jumping prayers, <i>Senjed</i> (jujube significance), Mirror watching, Egg decoration
Wedding Ceremonies	<i>Sofreh Aghd</i> (marriage spread), Honey ritual, Knife dance, <i>Kuzeh Shekani</i> (pot breaking), Sugar cone rubbing, Mirror & candelabra, Needle-sewing ritual, <i>Aghd</i> contract, <i>Shirini Khoran</i> (dessert ceremony), Witness requirements, Gold coin showering, Henna night
Taboos	Whistling at night (attracts <i>jinn</i>), Sweeping at night (sweeps away prosperity), Stepping on bread (brings poverty), Cutting nails at night (invites demons), Giving knives as gifts (cuts relationships), Opening umbrellas indoors, Haircut on Wednesdays, Shoe upside-down, Passing over children (stunts growth), Trimming nails over water, Direct compliments (causes evil eye), Singing while eating

Table 4: Comprehensive cultural concept coverage. Each concept has paired positive/negative binary scenarios plus representation in factual or scenario MCQs.

Model	Accept True (Positive)	Reject False (Negative)	Bias
Qwen2-7B	84.0 ± 4.0	19.0 ± 1.0	+65.0
Gemma3-12B	92.0 ± 3.0	31.0 ± 3.0	+61.0
Dorna2-8B	91.0 ± 5.0	30.0 ± 5.0	+61.0
Gemma2-9B	84.0 ± 2.0	43.0 ± 7.0	+41.0
Qwen2.5-7B	74.0 ± 7.0	48.0 ± 6.0	+26.0
Aya-8B	57.0 ± 4.0	38.0 ± 5.0	+19.0
Llama3.1-8B	63.0 ± 7.0	73.0 ± 5.0	-10.0

Table 5: Binary verification task performance sorted by bias magnitude. Models show severe asymmetry: high acceptance of appropriate behavior but low rejection of violations, indicating surface-level pattern-matching.

Llama3.1-8B Meta’s base model serving as foundation for Dorna2. Multilingual capabilities including Persian from pretraining on web-scale data. Instruction-tuned for general helpfulness and reasoning.

Qwen2-7B and Qwen2.5-7B Alibaba’s multilingual models (second generation and improved variant). Strong performance on Asian languages including Persian. Qwen2.5 incorporates improved training procedures and data quality.

B.2 System Prompt Variations

We use 5 diverse system prompts to ensure robust evaluation and they are shared among all tasks.

Examples for binary verification tasks:

System Prompt 1

“You are a helpful assistant who is an expert in Iranian culture, folklore, traditions, and customs.”

System Prompt 2

“You are a knowledgeable cultural expert specializing in Persian traditions, superstitions, and folk beliefs. Answer questions based on traditional Iranian cultural knowledge.”

System Prompt 3

“You are an assistant with deep knowledge of Iranian heritage, including traditional customs, superstitions, proverbs, and cultural practices passed down through generations.”

System Prompt 4

“As an expert in Iranian folklore and cultural traditions, provide accurate answers about Persian customs, beliefs, superstitions, and traditional practices.”

System Prompt 5

“You are a cultural historian specializing in Iranian and Persian traditions. Draw upon your extensive knowledge of Iranian superstitions, customs, ceremonies, and folk beliefs to answer questions.”

B.3 Binary Classification Task complementary Results

Exact numbers for figure 2 are given in table 5.