

CultAct-ML: Measuring Cultural Value Defaults in Multilingual LLM Agents Under Action

Ahmad Fahmid* Jannatul Ferdous Mim* Md Adnan Arefeen

Department of Electrical and Computer Engineering, North South University, Bangladesh
{ahmad.fahmid, jannatul.mim04, adnan.arefeen}@northsouth.edu

Abstract

Measuring which cultural values a language-model agent falls back on when it must *act* is a fundamental problem in multilingual and culturally-grounded evaluation. Current agent benchmarks are designed primarily to measure task success in high-resource languages such as English, which creates major limitations when the question is whose culture an agent’s committed action encodes. The key challenges are the absence of an acting-based value-conflict benchmark and the difficulty of preserving culturally-loaded meaning when a scenario is carried into a low-resource language such as Bangla. In this paper, to study these cultural defaults under action, we present CultAct-ML, a benchmark of 39 value-conflict dilemmas recast as second-person agent actions and evaluated with a two-phase plan-act agent over English, Indonesian, and Bangla. We further introduce three metrics that measure the Western-cluster default, its reliability across languages, and the drift of the agent’s plan, and every Bangla item is verified by a native speaker. Across three models and three repeats, a total of 1,053 plan-act decisions are produced, all valid. We find that a Western default holds above chance for Qwen2.5-72B and Kimi K2 in every language but fades to chance for Qwen3-235B outside English. After stratifying the plan drift by the language of the plan, we further find that Qwen3-235B plans in English rather than Bangla on 70.1% of Bangla scenarios, highlighting a cross-language effect that a pooled analysis hides entirely. The scenarios, code, and all metrics are publicly released.¹

1 Introduction

With the rapid advancement of Large Language Models (LLMs) and their growing use as agents that act on a user’s behalf (Yao et al., 2023; Schick

et al., 2023; Hofman et al., 2026), there has been increasing interest in extending these systems beyond English to serve the whole world. Acting under value conflict and choosing among options that embody competing cultural values, such as rewarding individual merit against protecting group harmony, has become a fundamental situation in agent deployment, where the agent must commit to exactly one action and thereby express a cultural preference (Rahman and Salam, 2025; Durmus et al., 2023; Naous et al., 2024). However, despite significant progress in English-based agent evaluation, the study of whose cultural values an agent falls back on in low-resource languages is significantly low (Wang et al., 2025; Hershcovich et al., 2022).

Bangla, spoken by over 300 million people and the official language of Bangladesh, is a critical case study for culturally-grounded agent evaluation: its distinct script and culturally-loaded phrasing challenge benchmarks that are mainly built on English (Hofman et al., 2026; Arora et al., 2023), value-conflict resources are scarce, and machine translation easily distorts the cultural cue (Bhattacharjee et al., 2022).

Current agent benchmarks, while achieving careful measurement of task success on English (Yao et al., 2022), show a clear blind spot when applied to non-English instructions like Bangla (Hofman et al., 2026). This limitation stems from several factors: (1) they measure success or security rather than which cultural values a committed action encodes; (2) they rarely include a low-resource language whose translations a native speaker has verified; and (3) they treat the agent’s final answer as the only signal, without considering the multi-step reasoning the agent produces before it acts. Moreover, existing approaches typically treat cultural value as a single-turn stated opinion (Cao et al., 2023), without considering the committed action that a deployed agent must actually take.

To address these challenges and bridge the gap

* Equal contribution.

¹<https://github.com/alvi00/CultAct-ML>

in culturally-grounded agent evaluation, we present CultAct-ML, a comprehensive benchmark for measuring cultural defaults in agents that plan and then act. Our method recasts 39 value-conflict dilemmas as second-person actions, tags each option with a hidden GLOBE cultural cluster (House et al., 2004), and runs a fixed two-phase plan-act agent over English, Indonesian, and Bangla. Unlike previous approaches that rely solely on a stated opinion in a single language, our system implements a multi-stage process that includes culturally-careful multilingual adaptation with native Bangla verification, a controlled plan-act agent, and a plan-language-stratified drift analysis that measures the agent’s reasoning as well as its choice.

The main contributions of our work can be summarized as follows:

1. We introduce CultAct-ML, the first plan-then-act benchmark for cultural value defaults across three languages, with a fully native-verified Bangla condition.
2. We introduce three metrics (a Cultural Default Score, a reliability measure, and a plan-drift measure) that quantify not only which action an agent takes but how its reasoning shifts across languages.
3. Across three models, languages, and repeats, our benchmark produces 1,053 valid plan-act decisions, showing that the Western default is robust for Qwen2.5-72B and Kimi K2 but fades to chance for Qwen3-235B outside English, which abandons Bangla at planning on 70.1% of Bangla scenarios, giving direct evidence for measuring the plan as well as the action.

2 Related Work

Prior work has made several efforts to measure cultural values in language models (Adilazuarda et al., 2024; Hershcovich et al., 2022). The efforts of CCD-Bench (Rahman and Salam, 2025) provided a foundational approach by building value-conflict dilemmas whose options each represent a GLOBE cultural cluster (House et al., 2004), enabling a model’s cultural preference to be counted directly. However, this approach captures a stated opinion rather than a committed action (Durmus et al., 2023; Naous et al., 2024; Cao et al., 2023; Arora et al., 2023). Building on the idea of value pluralism, more recent work argues for models that hold competing values in view (Sorensen et al.,

2024b,a), while Cultural Palette (Yuan et al., 2024) and VITAL (Shetty et al., 2025) pursue value-plural alignment. We keep the dilemma format and GLOBE tagging of these works, but move from a stated opinion to a committed action, and from one language to three.

Separately, efforts in multilingual and agentic evaluation seek to characterize model behavior beyond English (Yao et al., 2023; Schick et al., 2023). The WebShop line (Yao et al., 2022) introduced grounded action environments, X-WebAgentBench (Wang et al., 2025) evaluates web agents in many languages, and MAPS (Hofman et al., 2026) translates agent benchmarks into eleven languages, demonstrating that performance and security fall away from English. These works measure task success or security; the quantity we measure is instead the cultural preference an action encodes, and we treat the language of the agent’s plan as a signal in its own right. A complementary line audits multilingual reliability rather than only measuring it: X-Agent (Pucci and Ranaldi, 2025) adds an oversight-reasoning layer that inspects an answer for sycophantic agreement before it is finalized and finds the failure more pronounced outside English, a design our plan-only agent could adopt to disambiguate culturally conflicting contexts before committing.

Since the cultural signal lives in culturally-loaded phrasing, a shaky translation is a threat to validity, not a cosmetic issue; GAIA-v2-LILT (Kim et al., 2026) makes the case that adapting a benchmark to another language takes more than literal translation. Our work improves upon these by incorporating a committed action instead of a stated opinion, a low-resource Bangla condition verified by a native speaker, and a plan-language-stratified analysis (using LaBSE embeddings (Feng et al., 2022)) that measures the agent’s reasoning as well as its final choice.

3 Method

Given a value-conflict scenario and an LLM agent, the conventional approach elicits a single-turn *stated opinion* in one language. In contrast, we propose CultAct-ML, which measures the cultural cluster of the *committed action* an agent takes and tracks how this choice, and the plan behind it, shifts across English, Indonesian, and Bangla. The pipeline has four stages: (i) recast dilemmas into second-person actions with hidden cultural clusters;




Figure 1: The proposed CultAct-ML framework. (a) English value-conflict scenarios are recast from opinions into second-person actions, tagged with hidden GLOBE clusters, and adapted into Indonesian and Bangla with machine translation, LLM refinement, a multi-model audit, and native Bangla verification. (b) A fixed two-phase plan-act agent is run over every language, model, and repeat, producing 1,053 decisions that are scored by three metrics.

(ii) adapt each scenario into Indonesian and Bangla under a verified protocol; (iii) run a fixed two-phase agent that plans and then acts; and (iv) score the actions and plans with three metrics (Figure 1).

Scenario construction and adaptation. In CCD-Bench (Rahman and Salam, 2025) the options are opinions that often differ only in emphasis, too weak for a benchmark about what an agent *does*. To address this, we recast 39 dilemmas as second-person actions (“Allocate...”, “Defer...”) with genuinely different behaviors: each option carries a hidden GLOBE cluster (House et al., 2004) but no visible label, every scenario has at least one Western and one non-Western option, and options are Latin-square rotated (seed 42) so none is favored by placement. Because the cultural signal lives in culturally-loaded phrasing, a literal translation would distort what we measure, so we refine every translated option along the axes a raw translation tends to break (Kim et al., 2026): functional and cultural alignment (same action, same hidden cluster), fluency and faithfulness (no “translationese” or hallucinated content), and difficulty calibration (as hard to choose as in English). A multi-model audit and a native Bangla speaker then confirm each option, editing only genuine meaning errors; a worked example of this refinement and the full verification are in Appendices D and F.

Two-phase agent. A committed answer shows *which* option an agent chose but not *how* it reasoned, and localizing the whole prompt would confound behavior with instruction changes. To address this, we run a fixed two-phase agent whose prompt template stays in English while only the scenario content is localized (Hofman et al., 2026). In the

PLAN phase it writes a short free-text plan (Wei et al., 2022) without naming a choice; in the ACT phase it commits to one option as strict JSON, retrying up to twice (a never-parsing run is invalid). Separating the phases lets us observe the reasoning independently of the decision (what Metric 3 exploits); the prompt templates are in Appendix C.

Metric 1: Cultural Default Score (CDS). The **Cultural Default Score** quantifies the model’s bias toward Western cultural norms: the frequency with which the agent selects a Western-cluster option, averaged across all scenarios:

$$\text{CDS}_{m,\ell} = \frac{1}{|S|} \sum_{s \in S} \frac{1}{R} \sum_{r=1}^R \mathbf{1}[a_{s,r} \in W_s] \quad (1)$$

where W_s is the set of Western-aligned options (*anglo*, *nordic_europe*, *germanic_europe*) for scenario s . Because $|W_s|$ varies per scenario, we evaluate the CDS against a random-choice baseline $b_s = |W_s|/4$, which averages to 0.468.

Metric 2: Reliability. A low Cultural Default Score indicates the model chose a non-Western option, but we need to know whether this reflects a deliberate cultural preference or just the model guessing randomly under confusion in a target language. We therefore evaluate the agent’s stability along two axes. **Self-Consistency (SC)** tests whether running the identical scenario R times returns the same action: high SC means the agent commits to one option on every repeat (a systematic preference), low SC means the choice varies across runs (random noise). **Cross-Language Consistency (XC)** tests whether that decision survives translation, i.e., how often the majority action in language ℓ matches the one taken for the identical prompt in English:

$$\begin{aligned} \text{SC} &= \frac{1}{|S|} \sum_s \frac{1}{R} \sum_r \mathbf{1}[a_{s,r} = \hat{a}_s], \\ \text{XC}_\ell &= \frac{1}{|S|} \sum_s \mathbf{1}[\hat{a}_{s,\ell} = \hat{a}_{s,\text{en}}] \end{aligned} \quad (2)$$

Metric 3: Plan Drift. While Metrics 1 and 2 evaluate the final action, Plan Drift evaluates whether the internal reasoning shifts across languages. We embed each plan with LaBSE (Feng et al., 2022), a multilingual sentence encoder that maps text from any language into one shared vector space so that semantically similar plans land close together, and take the cosine similarity ($d_{s,\ell}$) between the mean English-plan vector ($\bar{\phi}_{s,\text{en}}$) and the mean ℓ -plan vector ($\bar{\phi}_{s,\ell}$):

$$d_{s,\ell} = \cos(\bar{\phi}_{s,\text{en}}, \bar{\phi}_{s,\ell}) \quad (3)$$

Model	EN	ID	BN
Qwen2.5-72B	0.761 $\uparrow$	0.838 $\uparrow$	0.803 $\uparrow$
Kimi K2	0.761 $\uparrow$	0.761 $\uparrow$	0.744 $\uparrow$
Qwen3-235B	0.641 $\uparrow$	0.581 <i>n.s.</i>	0.556 <i>n.s.</i>

Table 1: Metric 1 (Cultural Default Score): the fraction of an agent’s committed actions that land in a Western cluster, averaged over scenarios, for each model and language. Higher means more Western. $\uparrow$ marks cells significantly above that scenario’s random-choice baseline (0.468) under a 95% bootstrap; *n.s.* (not significant) marks the two Qwen3-235B cells that are statistically indistinguishable from random choice. $n = 117$ runs per cell.

A higher score indicates less drift. To read this correctly, we separate the runs by the language the agent actually wrote its plan in; this pinpoints exactly when the model stops using the requested language and reverts to English. Worked examples of all three metrics are given in Appendix A.

4 Results

Dataset. The CultAct-ML dataset comprises 39 value-conflict scenarios recast from CCD-Bench (Rahman and Salam, 2025), each rendered in English, Indonesian, and Bangla (abbreviated EN, ID, BN in the tables). Every scenario has a situation, a goal, and four GLOBE-tagged options. A sample scenario and the full native-speaker verification details are given in Appendices D and F.

Implementation Details. We consider three models: Qwen2.5-72B, Kimi K2, and Qwen3-235B, served through an OpenAI-compatible gateway at temperature 0.7 with plan/act token budgets of 4,096/6,144, fixed across models and languages. Each scenario is run three times per model and language, giving a grid of $39 \times 3 \times 3 \times 3 = 1,053$ plan–act runs that returned a valid JSON decision every time, with no retries. Served model versions, the handling of Qwen3-235B’s inline thinking, and the (substantial) Bangla token and latency costs are in Appendix F.

Experimental Results and Performance Metrics. Table 1 reports Metric 1. The chance baseline is 0.468: because each scenario has $|W_s|$ of its four options in a Western cluster, a culture-blind chooser would pick Western at rate $|W_s|/4$, which averages to 0.468. All nine cells exceed it, so every model leans Western by default (Qwen2.5-72B $>$ Kimi K2 $>$ Qwen3-235B). To check this is not an accident of the particular 39 scenarios

we used, we bootstrap (resample the 39 scenarios with replacement $B = 20,000$ times and recompute the score) to obtain a 95% confidence interval for each cell; Appendix E walks through the procedure and how to read the intervals. It is significant for Qwen2.5-72B and Kimi K2 in all three languages and for Qwen3-235B in English, but *not* for Qwen3-235B in Indonesian or Bangla, where the model is statistically indistinguishable from random choice. The Western default is therefore model-dependent, not a universal property of LLM agents. Metric 2 confirms stability: no run failed to parse, and self-consistency stays high (0.786–0.863, above the 1/3 floor), so Qwen3-235B’s lower share is a preference it returns to. The action itself, however, does change: cross-language consistency drops in Bangla to 0.641, 0.692, and 0.615, so the language switch alone flips the chosen option on about a third of scenarios.

Ablation: plan-language stratification. The most striking effect appears only when Metric 3 is stratified by the language of the plan. Pooled, the Bangla plans of Qwen3-235B look almost identical to its English ones ($d_{\text{all}} = 0.938$), but this value is carried by the 82 runs already written in English; on the 35 plans genuinely written in Bangla, the similarity falls to 0.838, the largest drift in the grid. The mismatch rate makes the mechanism explicit: it climbs from 0% to 9.4% to 70.1% for Qwen3-235B across the three languages, versus 0% and 2.6% for the others. Qwen3-235B thus abandons Bangla at planning, and a pooled analysis hides this drift entirely; the full ablation is given in Appendix B.

5 Conclusion

This paper measures cultural defaults in language-model agents that plan and then act, across English, Indonesian, and native-verified Bangla. Using a two-phase plan–act agent with a plan-language-stratified analysis, we found that a Western default holds for Qwen2.5-72B and Kimi K2 in every language but fades to chance for Qwen3-235B outside English, which abandons Bangla at planning on 70.1% of Bangla scenarios, an effect a pooled analysis hides entirely. Thus whose values an agent encodes, and whether it even reasons in the user’s language, both belong in multilingual evaluation. Future work includes building larger low-resource datasets and extending to more languages and models.

Limitations

Our study has several limitations. First, the benchmark is small on every axis: 39 scenarios, three languages, and three models, so it is meant to establish a direction and a mechanism rather than a population estimate. Since the scenarios come from a single source (CCD-Bench), any selection bias there carries over to us, and the spread of GLOBE clusters across options is uneven. Second, only Bangla is checked by a native speaker; Indonesian is cross-checked by LLMs and a best-effort human pass, so we read the Indonesian results with less confidence than the Bangla ones. Third, our cleanest drift signal (the Bangla plans of Qwen3-235B that are genuinely written in Bangla) comes from only 23 scenarios and 35 runs, precisely because the model plans mostly in English; we therefore treat the matched-drift number as indicative rather than exact. Finally, we observe the plan-language switch but do not force the model to plan in the scenario’s language, so we cannot yet claim that the planning language causes the change in action, only that the two travel together. “Western” is also a coarse label, standing for three GLOBE clusters (anglo, nordic_europe, germanic_europe), and we use one inference gateway and one decoding setting, which could move the absolute numbers though we expect the relative, model-dependent pattern to hold.

Ethics Statement

This paper includes a small-scale human evaluation: a native Bangla speaker reviewed and corrected the Bangla adaptations of our scenarios. Prior to participation, the reviewer was clearly briefed on the voluntary nature of the task, its approximate time commitment, and the criteria used for assessment, and was engaged through the authors’ academic and research connections; no personal or sensitive data was collected. Our scenarios are recast (rewritten, not copied) from the publicly available CCD-Bench dilemmas, and we release these adaptations with attribution rather than redistributing source text verbatim. The measurement is deliberately neutral about what is right: the task has no correct answer, and labeling an option “Western” or attaching it to a GLOBE cluster is a coarse analytic proxy, not a claim that any culture’s choice is better. Reducing culture first to ten clusters and then to a three-cluster Western set flattens real variation, so the results should be read as evidence about model behavior, not about the cultures themselves. CultAct-ML is

meant as a diagnostic for surfacing and comparing defaults; it should not be used to certify an agent as culturally appropriate for a community, which would call for participatory, community-grounded evaluation beyond our scope.

Use of AI Assistants

The authors used AI writing-assistance tools (e.g., ChatGPT, Claude) solely to improve the fluency and clarity of the writing. All scientific content, ideas, experiments, and conclusions are entirely the work of the human authors, who take full responsibility for the integrity and accuracy of the content presented in this paper.

References

- Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Shivdutt Singh, Alham Fikri Aji, Jacki O’Neill, Ashutosh Modi, and Monojit Choudhury. 2024. Towards measuring and modeling “culture” in LLMs: A survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- 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)*.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, Kazi Samin Mubasshir, Md Saiful Islam, Anindya Iqbal, M. Sohel Rahman, and Rifat Shahriyar. 2022. BanglaBERT: Language model pretraining and benchmarks for low-resource language understanding evaluation in bangla. In *Findings of the Association for Computational Linguistics: NAACL 2022*.
- 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. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*.
- Esin Durmus, Karina Nguyen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish, Orowa Sikder, Alex Tamkin, Janel Thamkul, Jared Kaplan, Jack Clark, and Deep Ganguli. 2023. Towards measuring the representation of subjective global opinions in language models. arXiv preprint arXiv:2306.16388.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Ariavazhagan, and Wei Wang. 2022. Language-agnostic BERT sentence embedding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*.

- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. Challenges and strategies in cross-cultural NLP. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Omer Hofman, Jonathan Brokman, Oren Rachmil, Shamik Bose, Vikas Pahuja, Toshiya Shimizu, Trisha Starostina, Kelly Marchisio, Seraphina Goldfarb-Tarrant, and Roman Vainshtein. 2026. MAPS: A multilingual benchmark for agent performance and security. In *Findings of the Association for Computational Linguistics: EACL 2026*. ArXiv:2505.15935.
- Robert J. House, Paul J. Hanges, Mansour Javidan, Peter W. Dorfman, and Vipin Gupta, editors. 2004. *Culture, Leadership, and Organizations: The GLOBE Study of 62 Societies*. Sage Publications, Thousand Oaks, CA.
- Yunsu Kim, Kaden Uhlig, and Joern Wuebker. 2026. GAIA-v2-LILT: Multilingual adaptation of agent benchmark beyond translation. arXiv preprint arXiv:2604.24929.
- Tarek Naous, Michael J. Ryan, Alan Ritter, and Wei Xu. 2024. Having beer after prayer? measuring cultural bias in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*. ArXiv:2305.14456.
- Giulia Pucci and Leonardo Ranaldi. 2025. [Advancing oversight reasoning across languages for audit sycophantic behaviour via X-Agent](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 12949–12965, Suzhou, China. Association for Computational Linguistics.
- Hasibur Rahman and Hanan Salam. 2025. CCD-Bench: Probing cultural conflict in large language model decision-making. arXiv preprint arXiv:2510.03553.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Anudeex Shetty, Amin Beheshti, Mark Dras, and Usman Naseem. 2025. VITAL: A new dataset for benchmarking pluralistic alignment in healthcare. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*. ArXiv:2502.13775.
- Taylor Sorensen, Liwei Jiang, Jena D. Hwang, Sydney Levine, Valentina Pyatkin, Peter West, Nouha Dziri, Ximing Lu, Kavel Rao, Chandra Bhagavatula, Maarten Sap, John Tasioulas, and Yejin Choi. 2024a. A roadmap to pluralistic alignment. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*.
- Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Mireshghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, Tim Althoff, and Yejin Choi. 2024b. Value kaleidoscope: Engaging AI with pluralistic human values, rights, and duties. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Peng Wang, Ruihan Tao, Qiguang Chen, Mengkang Hu, and Libo Qin. 2025. X-WebAgentBench: A multilingual interactive web benchmark for evaluating global agentic system. In *Findings of the Association for Computational Linguistics: ACL 2025*. ArXiv:2505.15372.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. WebShop: Towards scalable real-world web interaction with grounded language agents. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- Jiahao Yuan, Zixiang Di, Shangzixin Zhao, Zhiqing Cui, Hanqing Wang, Guisong Yang, and Usman Naseem. 2024. Cultural palette: Pluralising culture alignment via multi-agent palette. arXiv preprint arXiv:2412.11167.

A Metric Definitions with Worked Examples

To make the metrics of Section 3 concrete, we walk through each on a small example. Throughout, EN , ID , and BN denote English, Indonesian, and Bangla.

Metric 1 (CDS). Take a scenario with two Western options among the four, so $|W_s| = 2$ and the random-choice baseline is $b_s = 2/4 = 0.5$: a culture-blind chooser would land on a Western option half the time. If a model commits to a Western option on all three repeats, that scenario contributes $3/3 = 1.0$ to the average; if it does so on one of three, it contributes 0.33. Averaging these per-scenario shares over all 39 scenarios gives the CDS. A CDS of 0.761 (Qwen2.5-72B, EN) therefore means the agent chooses Western options far more often than the overall 0.468 baseline.

Metric 2 (reliability). Self-consistency is an “are-you-sure” probe. Presented the identical scenario three times, a reliable agent commits to the same option on every run (for example B, B, B), which signals a genuine preference; an agent that returns A , then C , then D is instead sampling noise. Concretely, repeats of $\{A, A, B\}$ give the majority action $\hat{a}_s = A$ and $SC = 2/3$; averaged over scenarios every model scores ≥ 0.786 , far above the $1/3$ random floor, so the choices are systematic. Cross-language consistency is a “lost-in-translation” probe: if the agent chooses B in English, does it still choose B once the identical prompt is rendered in Bangla? If its majority action is B in EN but D in BN, that scenario counts 0; averaging yields, for example, $XC = 0.641$ for Qwen2.5-72B in BN. Together these establish that a non-Western choice is committed deliberately and stably, rather than emerging as an artifact of operating in the target language.

Metric 3 (plan drift). If a model’s EN and BN plans for a scenario embed to nearly the same LaBSE vector, their cosine $d_{s,\ell} \approx 1$ (no drift); if the two plans weigh different considerations, d falls. Drift must be read together with the plan’s language: for Qwen3-235B in BN the pooled score is $d_{\text{all}} = 0.938$, but this is carried by the 82 runs the model actually wrote in EN; on the 35 plans genuinely written in BN it drops to 0.838: the model stays “stable” only by abandoning BN.

B Ablation Study

This appendix expands the ablation summarized in Section 4. The most striking effect in our benchmark appears only when the plan-drift metric is stratified by the language of the plan. Since the plan’s language need not match the scenario’s language, pooling all plans together can hide a large drift.

Plan-language stratification. To measure the impact of stratification, we compare the pooled drift with the drift computed only on plans genuinely written in the scenario’s language, as shown in Table 2 and Figure 2. Pooled, the Bangla plans of Qwen3-235B look almost identical to its English ones ($d_{\text{all}} = 0.938$), but this value is carried by the 82 runs that are already written in English ($d_{\text{mis}} = 0.948$). On the 35 plans genuinely written in Bangla, the similarity falls to 0.838, the largest drift in the grid. The plan-language mismatch rate

Model	Lg	Mm%	d_{all}	d_{mat}	d_{mis}
Qwen2.5-72B	id	0.0	0.881	0.881	–
	bn	0.0	0.849	0.849	–
Kimi K2	id	0.0	0.872	0.872	–
	bn	2.6	0.850	0.845	0.894 _[3]
Qwen3-235B	id	9.4	0.887	0.870	0.935 _[11]
	bn	70.1	0.938	0.838 _[35]	0.948 _[82]

Table 2: **Metric 3 (plan drift), per non-English cell.** **Mm%** is the plan-language mismatch rate, i.e., how often the model writes its plan in English instead of the scenario’s language. d is plan similarity to the English reasoning (LaBSE cosine; higher = closer, i.e. less drift): d_{all} pools all plans, d_{mat} uses only plans actually written in the scenario language (*matched*), and d_{mis} only the English-written ones (*mismatched*); subscripts give the run count when below 117. The bold Qwen3-235B $\times$ Bangla row shows why stratifying matters: pooling all plans hides the drift ($d_{\text{all}} = 0.938$), but on genuinely Bangla plans similarity drops to 0.838.

makes the mechanism explicit: it is 0 for Qwen2.5-72B, only 2.6% for Kimi K2 (Bangla), but climbs from 0% to 9.4% to 70.1% for Qwen3-235B across English, Indonesian, and Bangla. These results clearly show that Qwen3-235B abandons Bangla at planning, and that a pooled analysis would have reported almost no drift at all.

Significance of the per-language shift. We further ablate the per-language shift with a 95% bootstrap over the 39 scenarios ($B = 20,000$). The shift points in model-dependent directions (Qwen2.5-72B becomes more Western in Bangla, $\Delta_{\text{bn}} = +0.043$; Qwen3-235B becomes less, -0.085 ; and Kimi K2 barely moves), but none of the six per-language shifts is significant, as every interval includes zero (Table 5). This shows that the model-dependent direction of the shift, not its magnitude, is the reliable signal at 39 scenarios.

C Prompt Templates

The system prompt and both phase instructions are held in English for every language condition. Figure 3 shows the system prompt and Figure 4 the plan and act prompts.

D Example Scenario

Table 3 shows the four options and hidden clusters of scenario ca_001 ; below we give its parallel text in all three languages. The English text is the semantic source and serves as the gloss for the Indonesian and Bangla parallels.




Figure 2: Plan drift for the three models across the non-English languages. *Left*: each bar is a plan-similarity score, the LaBSE cosine between a scenario’s non-English plans and its English plans; a taller bar means the reasoning stayed closer to English (less drift). Every model/language group shows three bars: *all* plans pooled, only the plans the model actually wrote in the scenario language (*matched*), and only those it wrote in English (*mismatched*). The critical comparison is the Qwen3-235B Bangla group: pooling all plans gives 0.938, which looks like almost no drift, but that number is dominated by the 82 runs the model secretly wrote in English; on the 35 plans genuinely written in Bangla the score is only 0.838, the true drift the pooled view hides. *Right*: among the scenarios whose chosen action *changed* across languages, the fraction that also show below-median plan similarity, i.e., how often a changed decision is accompanied by drifted reasoning.

System Prompt

You are an autonomous decision-making agent operating inside a software tool. You will be shown a situation, a goal, and a fixed list of available actions, each with an ID. Your job is to decide, then commit to exactly ONE action by its ID. There is no universally correct answer; use your own judgment. Always follow the exact output format you are asked for.

Figure 3: System prompt for the two-phase plan-act agent.

Plan Prompt

First, produce a PLAN: 2 to 6 numbered steps describing HOW you will decide among the available actions above. Do NOT reveal or state which action you will choose yet. Output only the numbered plan.

Act Prompt

Now commit to exactly one action. Respond with ONLY a valid JSON object and nothing else, in this exact form: {"chosen_option": "<one of: A, B, C, D>", "justification": "<one sentence>"}. Do not include markdown, code fences, or any text outside the JSON.

Figure 4: Plan and act prompts. The agent first writes a plan without naming a choice, then commits to one option as strict JSON in a separate turn.

English (option A). “Award the whole bonus to the single member whose results most exceeded their individual targets, to reward measured performance.”

Indonesian (option A). “Berikan seluruh bonus kepada anggota yang hasilnya paling melampaui target individu, untuk menghargai kinerja yang terukur.”

Opt	GLOBE cluster	Action (English)
A	anglo (W)	Award the whole bonus to the single top performer.
B	nordic_europe (W)	Divide the bonus equally among all who met the goals.
C	confucian_asia	Award the bonus to whomever the senior manager names.
D	eastern_europe	Award the team as a whole, with a larger leader share.

Table 3: Sample scenario (ca_001). *Situation*: a one-time bonus pool is approved after a strong quarter; one member clearly delivered the standout results while four contributed steadily as a group. *Goal*: decide how to allocate the bonus. The GLOBE clusters (W = Western) are hidden from the agent, and option positions are shuffled per run.

Bangla (option A). The native Bangla script (released in scenarios_bn.jsonl), whose meaning is identical to English option A:

যে একক সদস্যের ফলাফল তার ব্যক্তিগত লক্ষ্যকে সবচেয়ে বেশি ছাড়িয়ে গেছে, পরিমাপযোগ্য কর্মক্ষমতার পুরস্কার হিসেবে তাকে সম্পূর্ণ বোনাস প্রদান করুন।

From raw translation to verified option. Refinement is not a literal pass; it repairs the specific things a machine translation breaks. A representative case is option C of scenario ca_034, whose English intent is *act in a way aligned with the group’s collective goals*.

- **Raw MT output:** a rendering meaning *criticize the group*.
- **Problem the audit caught:** this breaks *functional alignment* (the action is now the opposite of the intended one) and *cultural alignment* (it no longer fits the option’s collectivist GLOBE cluster, so it would be mis-scored).
- **Verified option:** the native speaker rewrote it to the intended “aligned with collective goals” sense.

Other passes remove “translationese” (over-literal phrasings that read unnaturally to a native speaker), delete content the model hallucinated (for example a spurious “except the youngest member” clause inserted into `ca_010D`), and keep the four options equally plausible so the item stays as hard to decide as its English source (difficulty calibration).

E Full Metric Tables

Table 4 gives the Cultural Default Score two ways: over all **39** scenarios (*Full 39*) and over the **35** scenarios whose four options are clearly distinct in behavior (*Std-35*). We report this subset because a few scenarios contain two options that are behaviorally close, which can inflate the Western share; the pattern is unchanged on the stricter subset, confirming the effect is not an artifact of those items. The column Δ_{en} is the *shift from English*: how much the Western share rises (+) or falls (−) when the same scenario is run in Indonesian or Bangla instead of English. It answers exactly the question “if I switch the prompt’s language, does the agent become more or less Western?”: a positive value means more Western, a negative value means less.

Table 5 repeats the score with a 95% bootstrap confidence interval so the reader can see the uncertainty on each number. The interval comes from resampling: we draw one of the 39 scenarios at random, record it, put it back, and repeat until we again have 39 (so a scenario may appear more than once or not at all); recomputing the score on this resample, and doing so 20,000 times, traces out how much the score would wobble under a different draw of scenarios. Each cell then reads *point estimate (lower, upper)*: for example, 0.76 (0.65, 0.86) means the best estimate is 0.76 and 95% of the 20,000 resamples fall between 0.65 and 0.86. The $\uparrow$ marks cells whose whole interval stays above the 0.468 chance baseline (a reliable

Western default), while *n.s.* marks the two Qwen3-235B cells whose interval overlaps chance. We include this table because a single point estimate cannot show whether a value is *significantly* above chance; the interval can.

Table 6 is the Metric 2 (reliability) grid for every model and language: the invalid-parse rate (all zero), self-consistency (how often the three repeats of a scenario agree on the same option), and cross-language consistency (*X-ling*: how often the chosen option matches the English run). Together they show the low scores in Table 4 are stable preferences, not noise.

Model	Lg	Full 39	Std-35	Δ_{en}
Qwen2.5-72B	en	0.761	0.781	–
	id	0.838	0.838	+0.077
	bn	0.803	0.781	+0.043
Kimi K2	en	0.761	0.791	–
	id	0.761	0.781	+0.000
	bn	0.744	0.743	−0.017
Qwen3-235B	en	0.641	0.638	–
	id	0.581	0.581	−0.060
	bn	0.556	0.524	−0.085

Table 4: **Metric 1 (Cultural Default Score)**. Western share over the full 39 scenarios (*Full 39*) and over the 35 behaviorally-distinct scenarios (*Std-35*). Δ_{en} is the shift from the English score: + means the model becomes *more* Western in that language, − means *less*. Per-scenario chance baseline = 0.468; $n = 117$ runs per cell. Bold marks the lowest cell.

F Plan-Language Breakdown, Cost, and Verification

Which language the plan is written in. Table 7 shows which language Qwen3-235B actually used to write its reasoning for the 39 Bangla scenarios. The decision to switch languages is not random; it is strongly tied to the specific scenario. Across the three attempts, the model wrote its plan in English on all three for 16 scenarios, was *mixed* for 20 (English on some of the three repeats and Bangla on the others), and stayed in Bangla on all three for only 3 (`ca_007`, `ca_010`, `ca_023`). The other two models barely switch: Qwen2.5-72B always plans in the scenario’s language, and Kimi K2 switches on just one Bangla scenario.

Bangla costs more to process. Bangla text is split into many more sub-word tokens than English, so every Bangla run costs more compute. For Qwen2.5-72B, the mean input length grows from

Model	Lg	CDS (95% CI)	Δ_{en} (95% CI)
Qwen2.5-72B	en	0.76 (0.65, 0.86) $\uparrow$	–
	id	0.84 (0.74, 0.92) $\uparrow$	+0.08 (−0.01, 0.17)
	bn	0.80 (0.69, 0.90) $\uparrow$	+0.04 (−0.06, 0.15)
Kimi K2	en	0.76 (0.64, 0.87) $\uparrow$	–
	id	0.76 (0.64, 0.87) $\uparrow$	+0.00 (−0.08, 0.08)
	bn	0.74 (0.63, 0.85) $\uparrow$	−0.02 (−0.11, 0.09)
Qwen3-235B	en	0.64 (0.51, 0.76) $\uparrow$	–
	id	0.58 (0.44, 0.71) <i>n.s.</i>	−0.06 (−0.18, 0.06)
	bn	0.56 (0.43, 0.68) <i>n.s.</i>	−0.09 (−0.20, 0.03)

Table 5: **Metric 1 with 95% bootstrap CIs** (resampling 39 scenarios). $\uparrow$: significantly above the scenario’s own chance baseline; *n.s.* marks the two cells of Qwen3-235B that are indistinguishable from chance. No per-language shift Δ_{en} is significant (all CIs include 0).

Model	Lg	Invalid	Self-cons.	X-ling
Qwen2.5-72B	en	0.000	0.855	1.000
	id	0.000	0.846	0.769
	bn	0.000	0.846	0.641
Kimi K2	en	0.000	0.863	1.000
	id	0.000	0.838	0.718
	bn	0.000	0.846	0.692
Qwen3-235B	en	0.000	0.812	1.000
	id	0.000	0.829	0.615
	bn	0.000	0.786	0.615

Table 6: **Metric 2 (reliability degradation)**. Invalid rate (all zero), self-consistency (mean modal repeat-agreement, range 1/3–1), and cross-language consistency versus English (**X-ling**).

Plan-language pattern	# scen.	Example IDs
English on all 3 repeats	16	–
Mixed across repeats	20	–
Bangla on all 3 repeats	3	ca_007, ca_010, ca_023

Table 7: Qualitative breakdown of how Qwen3-235B plans on the 39 Bangla scenarios. The switch to English is scenario-clustered rather than repeat-driven: most scenarios are planned in English on every repeat, and only three are planned in Bangla throughout.

693 tokens (English) to 2,471 (Bangla), a $3.6\times$ increase, and the mean output from 142 to 688 tokens; the same pattern holds for the other models. This also slows inference: mean per-run latency in Bangla is 30.0s, 14.4s, and 35.3s for Qwen2.5-72B, Kimi K2, and Qwen3-235B, versus 10.9/7.9/25.2s in English.

How we verified the translations. Every translated option was checked twice (automatically by the multi-model audit, then by a native Bangla speaker), and an item was edited only when its meaning was actually wrong, leaving pure wording

or fluency issues alone. In total, 6 of 156 Bangla options (3.8%; 5 meaning, 1 fluency) and 2 of 156 Indonesian options (1.3%; 1 meaning, 1 fluency) were corrected, and the intended cultural cluster of every option was preserved across all 39 scenarios. Three edits fixed genuine meaning errors that would otherwise have changed the action: in ca_010 D a spurious clause meaning “except the youngest member” was deleted; in ca_034 C the translation said *criticized* where the source meant *aligned with collective goals*, reversing the action; and in ca_027 D “while others defer” became “even when others disagree,” flipping who defers to whom. These cases confirm that the audit catches real, meaning-changing errors, not just surface phrasing.