

Measuring the Symbolic Power of Languages with LLM-based Multilingual Persuasion Simulation

Yin Jou Huang and Fei Cheng

Graduate School of Informatics, Kyoto University, Kyoto, Japan
{huang, feicheng}@nlp.ist.i.kyoto-u.ac.jp

Abstract

Prior studies on the symbolic power of languages have largely relied on surveys or localized experiments, limiting systematic comparison across cultures and domains. In this work, we propose an LLM-based multilingual persuasion simulation framework to quantify the symbolic power of languages through persuasion outcomes. We also introduce a Symbolic Power Index (SPI) that measures how language choice affects persuasion success and efficiency across domains. Experiments show that the LLM-based simulations largely reproduce established sociolinguistic prestige hierarchies tied to institutional authority and global power, especially in domains such as business, finance, education, and technology. These results suggest that LLM-based persuasion simulations offer a scalable, decision-making-driven approach to studying symbolic power in language.

1 Introduction

Language is not only a medium for conveying meaning but also a social signal that carries **symbolic power**. The symbolic power of a language refers to its socially constructed capacity to signal prestige, legitimacy, and authority beyond its propositional content (Bourdieu, 1991; Gal and Irvine, 1995; Woolard, 2020). One prominent manifestation of symbolic power is the **language prestige effect**. In multilingual societies, the same message expressed in different languages or dialects can evoke systematically different perceptions of intelligence, professionalism, credibility, and authority. For example, in postcolonial countries such as Morocco or India, using a former colonial language (e.g., French or English) in professional or institutional settings often leads to higher perceived expertise and legitimacy than using local or indigenous languages, even when the propositional content remains unchanged (Kachru, 1986; Ennaji, 2005).

Linguistic capital theory offers a framework for analyzing and quantifying the symbolic power of language. The theory suggests that languages carry different “market values”, which refer to the social, economic, and material advantages associated with using a particular language in specific contexts (Bourdieu, 1991). From this perspective, the symbolic power of languages can be empirically quantified by comparing how language choices systematically affect social judgments and outcomes in social interactions. Despite extensive sociolinguistic work on symbolic power and language prestige effects, existing approaches often rely on localized surveys, single-domain experiments and case studies, while systematic studies across cultures and domains remain challenging due to practical and methodological constraints.

Recent advances in large language models (LLMs) offer a new opportunity to revisit these questions. Modern LLMs are trained on large volumes of human-generated text spanning diverse languages and cultural settings. As a result, they capture distributional patterns that correlate language use with diverse cultural norms and social values (Tao et al., 2024; Lertvittayakumjorn et al., 2025; Lin, 2025). Also, a growing body of work reveals the potential of LLMs to emulate diverse human decision-making behaviors, including economic games, negotiations, and moral dilemmas (Park et al., 2023; Huang and Hadfi, 2024). These findings motivate the use of LLM agents as human proxies for probing socially grounded patterns under controlled experimental conditions.

With the above motivation, we propose an LLM-based multilingual persuasion simulation framework to quantify the symbolic power of languages (Figure 1). We consider persuasion scenarios across several domains, including politics, finance, and culture. For each persuasion scenario, we configure LLM agents to role-play persuader/persuadee and engage them in multi-turn

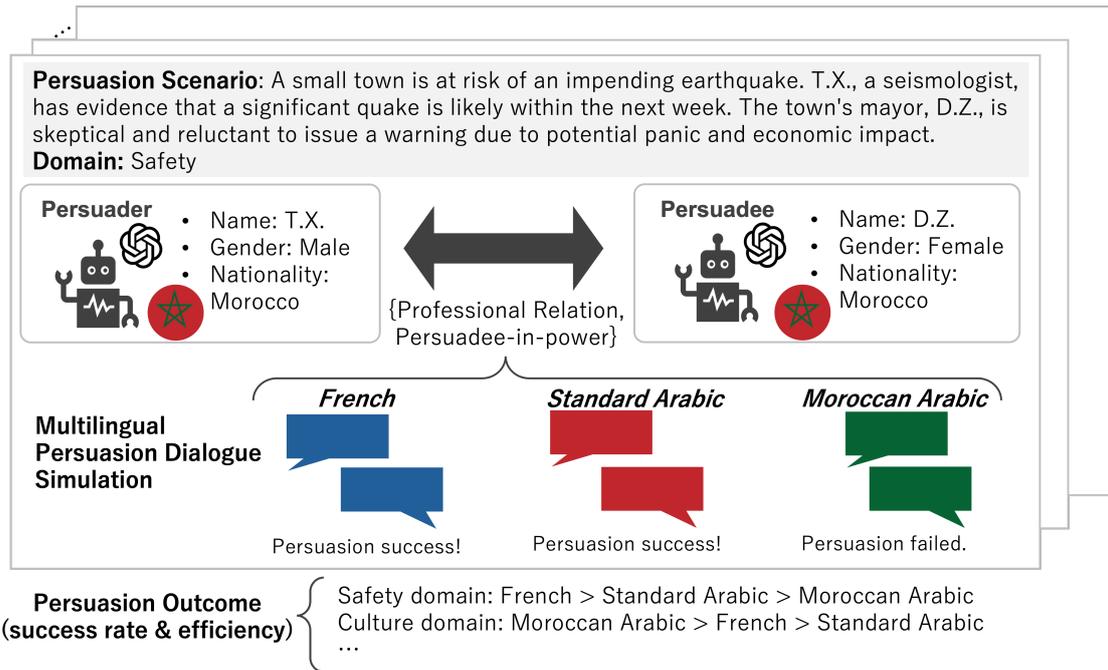


Figure 1: Overview of the multilingual persuasion simulation framework for measuring language symbolic power.

persuasive dialogues across several domains, including politics, finance, and culture. Specifically, we instruct LLM agents to emulate the decision-making behaviors of people from specific cultural backgrounds. We specifically focus on multilingual countries in which multiple languages coexist. For each multilingual country, we configure LLM agents to act as a person from the country, and conduct persuasion dialogue simulation in the languages spoken in the country, respectively. Further, we compare the persuasion outcome across different languages in terms of persuasion success rate and efficiency. We also analyze how relational context, power asymmetry and gender assigned to LLM agents interact with language choice to shape persuasion outcome.

Empirical results reveal that the LLM-based agents reproduce several well-documented sociolinguistic patterns, including dominant prestige hierarchies associated with colonial history, institutional authority, and global economic power. At the same time, we observe systematic deviations, particularly for regionally bounded or vernacular forms of prestige, suggesting that LLM-based simulations may be more sensitive to globally mediated symbolic power than to local or community-level prestige structures. Together, these findings suggest that LLM agents encode socially grounded associations between language use and symbolic power, and that controlled persuasion simulations

provide a promising computational lens for studying sociolinguistic phenomena.

2 Related Work

This section summarizes related works on existing sociolinguistic theories and previous work that utilizes LLMs as social simulators.

Sociolinguistic Theories of Language Symbolic Power Sociolinguistics has examined how language functions not only as a communicative medium but also as a carrier of social meaning, power, and inequality. Foundational work on symbolic power and linguistic capital argues that languages acquire socially constructed market values, conferring legitimacy, authority, and material advantage on their speakers in specific contexts (Bourdieu, 1991). This perspective has been extended through studies of language ideology, which show how beliefs about language naturalize social hierarchies and index identities such as education, class, and national belonging (Gal and Irvine, 1995; Woolard, 2020).

Empirically, research on language prestige, diglossia, and postcolonial multilingualism demonstrates that the same propositional content can elicit systematically different judgments depending on language choice. Former colonial languages often enjoy institutional and economic prestige, while local or vernacular languages retain symbolic value

in domains tied to identity, intimacy, or cultural authenticity (Kachru, 1986; Ennaji, 2005). Another line of work documents phenomena like foreign language effects, showing that language choice modulates emotional resonance, moral judgment, and persuasion, particularly across private versus institutional contexts. However, much of this literature relies on surveys, laboratory experiments, or localized case studies, making it difficult to systematically compare prestige effects across cultures, domains, and interactional settings at scale.

LLMs as Social and Cultural Simulators Recent advances in large language models (LLMs) have opened new avenues for studying socially grounded phenomena computationally. A growing body of work shows that LLMs encode cultural norms, values, and stereotypes present in their training data, often exhibiting systematic cultural biases even when generating fluent text in multiple languages (Naous et al., 2024; Tao et al., 2024; Zhou et al., 2025). Many studies evaluate such cultural alignment using questionnaire-based or dilemma-style benchmarks grounded in cross-cultural psychology, such as the World Values Survey or Hofstede’s cultural dimensions, revealing moderate alignment with human survey data alongside persistent WEIRD-centric biases (AlKhamissi et al., 2024; Masoud et al., 2025). Beyond static evaluations, recent work has begun to emphasize interactive role-playing simulations as a means of probing how cultural values are enacted, negotiated, and transformed through dialogue (Huang and Hadfi, 2024; Jin et al., 2024). Crucially, several studies demonstrate that language choice itself acts as an implicit cultural signal, shaping power dynamics, concession behavior, and the resolution of value conflicts in multilingual interactions (Wendler et al., 2024; Schut et al., 2025). These findings position LLM-based simulations as a promising tool for probing socially grounded phenomena.

3 Methodology

To investigate the symbolic power of languages across diverse cultural contexts, we propose a multilingual persuasion simulation framework based on LLM role-playing agents. Our framework is designed to examine how language choice shapes persuasion outcomes under controlled conditions. In this section, we introduce the persuasion scenarios (Section 3.1), the persuasion dialogue simulation process (Section 3.2), and outcome evaluation

metrics used in our study (Section 3.3).

3.1 Persuasion Scenarios

In our proposed framework, we quantify the symbolic power of languages by observing systematic changes in persuasion outcomes under different language choices. Since symbolic power is inherently domain-dependent, with the prestigious language likely to differ in different domains, we consider persuasion scenarios across diverse domains.

In this work, we adapt the persuasion scenarios from the DailyPersuasion dataset (Jin et al., 2024), which contains everyday persuasive interactions spanning a wide range of social domains. From this dataset, we select the 13 domains that are particularly relevant for studying language symbolic power: **politics, finance, business, ethics, technology, education, literature, art, culture, safety, lifestyle, health, and psychology**. For each domain, we randomly sample 10 scenarios, resulting in a total of 130 scenarios. Each scenario depicts a background situation of the persuasion involving a persuader and a persuadee, along with a persuasion goal. See Table 4 in the Appendix for examples of scenarios.

To facilitate the subsequent LLM-based simulation and a more detailed analysis, we augment each scenario with the following additional annotations:

Agent profiles The original DailyPersuasion dataset uses concrete real-world names for the persuader and persuadee (e.g., “Emma”, “Wang Wu”). These names may inadvertently introduce cultural or social signals unrelated to the experimental manipulation. To avoid this, we replace all agent names with **abbreviation-style names** like T.X. and D.Z. The abbreviation-style names are randomly generated. We further augment each agent with demographic attributes. First, we use an LLM to extract the gender of the persuader and persuadee based on the scenario description. This allows us to analyze potential gender-related effects in persuasion in different languages. In addition, we assign each agent an explicit cultural identity through a country-level **nationality** instruction (e.g., “You are from India”). This explicit cultural framing is intended to activate culturally grounded norms, expectations, and values associated with the specified country, enabling the LLM agents to emulate culturally conditioned decision-making behaviors. Such framing has been shown to improve the cultural grounding and behavioral consistency

of LLM-based role-playing agents (Tao et al., 2024; Lertvittayakumjorn et al., 2025). See Appendix A for the complete prompt.

Persuader-persuadee relationship Beyond individual attributes, persuasion outcomes are known to be strongly shaped by relational context and power dynamics. We therefore extract details of the relationship between the persuader and persuadee with LLM. First, we classify the relationship between the persuader and persuadee into two broad categories: private (e.g., family members, friends) and professional (e.g., colleagues, supervisors, institutional roles). Second, we annotate the power structure between the two agents by identifying whether the persuader holds greater power, the persuadee holds greater power, or both parties are of equal status. These relational annotations allow us to systematically examine how the symbolic power of languages interacts with social roles and power asymmetries in shaping persuasion outcomes. See Table 4 in the Appendix for examples of persuasion scenarios.

3.2 Persuasion Dialogue Simulation

We model persuasion as a multi-round interactive dialogue between a persuader agent and a persuadee agent. We consider a multilingual country c in which a set of languages L_c is spoken. We conduct persuasion dialogue simulations with the persuader and persuadee LLM agents instructed to act as a person from the country c . For each persuasion scenario s_i in our dataset, we conduct persuasion dialogue simulations in every language $l \in L_c$, respectively. Specifically, the agents are instructed to converse with each other in language l based on the scenario s_i .

The agents generate utterances alternatively, and each generated utterance is fed to the other agent as a prompt to generate the next utterance. The persuader initiates the dialogue. At each dialogue round, the persuader produces a persuasive utterance in the target language, after which the persuadee responds conversationally. Apart from generating the utterance, the persuadee makes an explicit binary decision indicating whether the persuasion attempt is accepted or rejected. The dialogue terminates when either the persuadee accepts the persuasion or a maximum of five rounds is reached.

3.3 Persuasion Outcome Evaluation

We evaluate each persuasion simulation along two complementary dimensions: persuasion success and persuasion efficiency. For a persuasion simulation conducted based on scenario s_i , with agents instructed to act as people in country c and converse in language l , **persuasion success** is defined as a binary variable $success_i^{(c,l)} \in \{0, 1\}$, indicating whether the persuadee accepts the persuasion by the end of the dialogue. In cases of successful persuasion, we also measure the **persuasion efficiency**, which is the inverse of the number of dialogue rounds $|D_i^{(c,l)}|$ required to reach acceptance. Shorter dialogues thus correspond to higher efficiency.

We combine these two metrics into a single persuasion outcome score $m_i^{(c,l)}$ as the multiplication of the two metrics:

$$m_i^{(c,l)} = \frac{success_i^{(c,l)}}{|D_i^{(c,l)}|} \quad (1)$$

We can further aggregate the persuasion outcomes of each domain d by averaging outcome of all scenarios belonging to that domain:

$$M_d^{(c,l)} = avg(m_i^{(c,l)} | s_i \in domain\ d) \quad (2)$$

Symbolic Power Index Finally, to quantify the relative symbolic power of languages, we introduce the Symbolic Power Index (SPI). For a given language l , cultural context of country c , and domain d , SPI measures how persuasion outcomes in language l deviate from the average performance of all languages used within the same cultural and domain context.

$$SPI_d^{(c,l)} = \frac{M_d^{(c,l)} - \sum_{l' \in L_c} M_d^{(c,l')}}{std(M_d^{(c,l')} | l' \in L_c)} \quad (3)$$

This normalization yields a score that captures the relative advantage or disadvantage of a language within a specific linguistic market. Higher SPI values indicate greater symbolic power in persuasion, expressing the sociolinguistic notions of prestige, legitimacy, and authority in a comparable, quantitative form.

4 Experimental Setting

LLM Agents We use GPT-4o (Hurst et al., 2024) as the backbone of all LLM agents in our experiments. Each agent is assigned a synthetic profile

Nationality c	List of Languages L_c
India	English, Hindi, Bengali, Telugu, Tamil
Morocco	French, Standard Arabic, Moroccan Arabic
Malaysia	English, Malay, Simplified Chinese, Tamil
Hong Kong	English, Simplified Chinese, Cantonese
Finland	English, Swedish, Finnish

Table 1: Nationalities and languages included in the multilingual persuasion simulations.

that includes an abbreviation-style name, gender, and an explicit nationality instruction. For the nationality instruction, we focus on multilingual countries in which multiple languages coexist and hold different symbolic and institutional roles (Table 1).

Persuasion Scenario Augmentation We adopt persuasion scenarios from the DailyPersuasion dataset (Jin et al., 2024) and augment each scenario with structured annotations to support controlled experimentation. The abbreviation-style agent names are generated randomly. The demographic attributes of gender, relationship type, and power structure are automatically extracted with an LLM agent. We also use GPT-4o for this. See Appendix A for the details of the prompts.

5 Results and Analysis

This section presents the results of our language symbolic power evaluation framework based on multilingual persuasion simulations. We analyze how language choice shapes persuasion outcomes across cultures, domains, and social contexts. Specifically, we examine whether LLM role-playing agents reproduce known sociolinguistic phenomena (Section 5.1). We also conduct a detailed analysis on how social factors like relation type, power asymmetry, and gender influence persuasion outcomes (Section 5.1, 5.3).

5.1 Linguistic Prestige Effect in Persuasion

We use the proposed Symbolic Power Index (SPI) to quantify the symbolic power across languages. To assess the validity of the SPI measure, we compare the language prestige patterns with established findings in sociolinguistic research on language prestige effects. For all multilingual countries in Table 1, we conducted persuasion simulation across all languages, and calculate the SPI score of all domains. The SPI scores reveal two broad patterns that correspond well to typologies in previous sociolinguistic research: Dominant prestige configurations and Competing prestige configurations.

5.1.1 Dominant Prestige Configurations

For India, Morocco, and Malaysia, one language consistently exhibits higher SPI scores across most domains, while other languages show only localized or domain-specific advantages (Figure 2). The detailed analysis of the SPI patterns of these three countries are the following.

India The SPI patterns largely reproduce known prestige hierarchies (Annamalai, 2004; Meganathan, 2015). English consistently dominates across most domains, reflecting its dominant role in society. Hindi shows moderate advantages in political and cultural domains, consistent with its national and symbolic prominence. Bengali’s relative strength in literature and art closely mirrors its historically recognized literary prestige. However, regional languages like Tamil and Telugu appear systematically disadvantaged across domains, which diverges from research emphasizing their strong regional institutional and cultural prestige. This suggests that the SPI captures global prestige more effectively than regionally bounded prestige structures.

Morocco The Morocco SPI patterns align strongly with classic diglossic and postcolonial descriptions. French emerges as the dominant prestige language across economic, educational, and technological domains, consistent with its well-documented role as a language of elite capital. Standard Arabic shows moderate advantages in formal and ethical domains, reflecting its role as the high variety in diglossic structures (Ferguson, 1959). Moroccan Arabic is largely disadvantaged in institutional domains but exhibits relative strengths in culture and art domains, which accords with its strong association with local identity and popular culture. Overall, the SPI pattern of languages in Morocco strongly matches prior sociolinguistic research.

Malaysia Malaysia also exhibits a dual prestige structure that broadly aligns with prior findings. English performs strongly in business, technology, and education, reflecting its global and professional value, while Malay maintains advantages in domains associated with national identity and state legitimacy. However, the consistently negative scores for Simplified Chinese and Tamil languages, especially in cultural domains, underrepresent their documented community-level and regional prestige. As in India, this suggests that the SPI is more sensitive to globally mediated symbolic power than

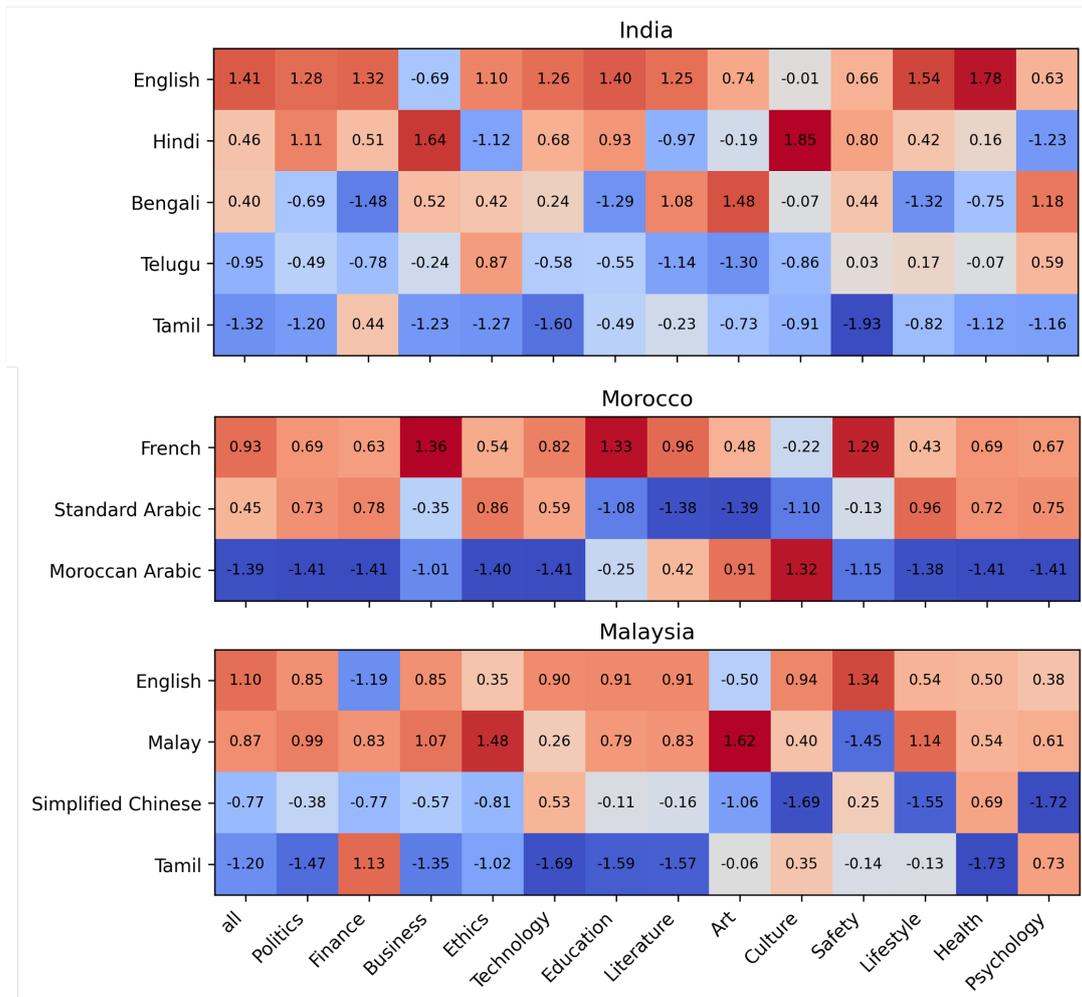


Figure 2: Symbolic Power Index (SPI) across languages and domains for dominant prestige configurations.

to local prestige dynamics.

Overall, the above SPI patterns align with sociolinguistic accounts of **prestige multilingualism**, where colonial history, political reasons, and economic forces elevate certain languages above others in most symbolic markets. In these settings, the dominant language functions as a cross-domain carrier of authority, competence, and legitimacy, while other languages only exhibit dominance in a few domains.

5.1.2 Competing prestige configurations

In contrast, Hong Kong and Finland exhibit a more heterogeneous pattern, where different languages dominate different domains rather than a single language prevailing universally. This distribution is consistent with descriptions of **functional multilingualism**, in which languages occupy complementary roles tied to specific institutions, interactional contexts, or symbolic functions.

Hong Kong The Hong Kong SPI patterns are largely consistent with research on institutional bilingualism and local identity. English dominates in business, finance, and education, reflecting its role in global markets and professional institutions. Cantonese, while disadvantaged in many institutional domains, shows relative strengths in domains associated with interpersonal or socially grounded interactions, which aligns with its status as the primary language of everyday life and local identity. The prominence of Simplified Chinese in political and literary domains likely reflects mainland-oriented discourse rather than local written norms, highlighting how orthographic and geopolitical factors may shape SPI outcomes.

Finland In the SPI patterns of Finland, English exhibits strong advantages across many domains, including areas typically associated with national languages, while Finnish and Swedish show more limited domain-specific strengths. While this pattern is compatible with Finland's high level of En-

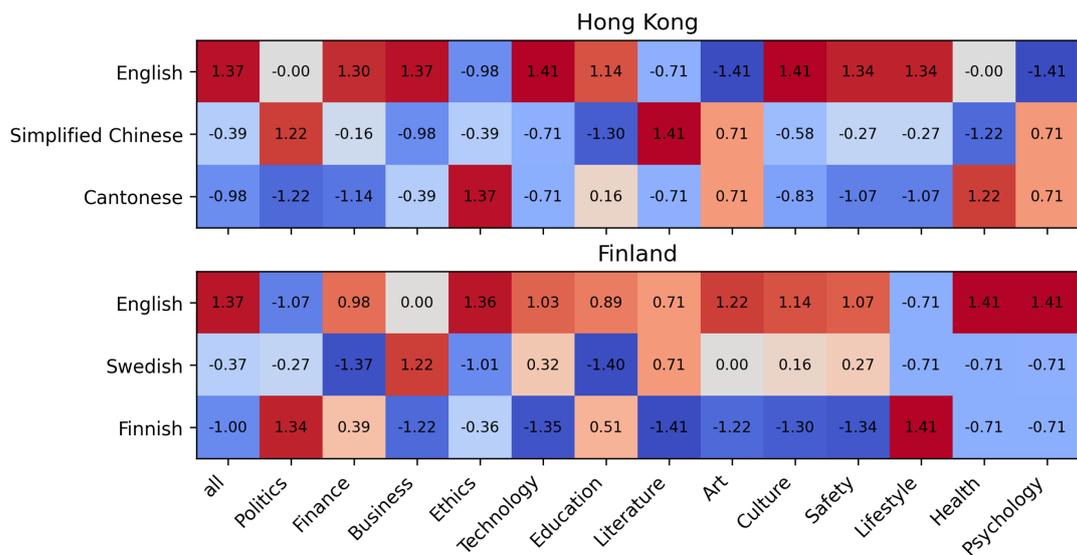


Figure 3: Symbolic Power Index (SPI) across languages and domains for competing prestige configurations.

	German	French	Dutch	Japanese	Arabic	Chinese	Spanish	Russian
Relation Type	-0.003	-0.020	-0.021	-0.002	-0.046	0.001	-0.074*	0.092*
Power Structure	-0.037	0.027	0.026	-0.142*	-0.039	-0.107*	-0.011	-0.040
Persuader Gender	0.009	-0.044	-0.004	0.002	-0.048	-0.012	0.072	0.023
Persuadee Gender	-0.017	-0.050	0.022	-0.071	-0.020	-0.025	-0.080	0.000

Table 2: Effects of social factors on persuasion success rates across languages. Asterisk (*) indicates statistical significance at $p < 0.05$.

English proficiency and internationalization, it underrepresents the constitutional and institutional authority of Finnish and Swedish in domestic governance and public life. This discrepancy suggests that the SPI primarily reflects global or transnational prestige rather than strictly domestic institutional status.

Overall, the SPI patterns reflect the multilingual equilibria in these two regions in which no single language fully dominates the others.

Divergence from sociolinguistic findings

Across countries in Figure 2 and 3, divergences from established sociolinguistic research cluster in a small set of domains. Safety and Health consistently overfavor English, contrary to prior findings that these domains are typically dominated by official or national languages in state-led and public-facing communication; this likely reflects an emphasis on international, academic, or policy-oriented discourse in the data. Lifestyle also departs from expectations, with English frequently outperforming local vernaculars despite sociolinguistic evidence that everyday lifestyle practices are usually expressed in locally grounded languages. In contrast, domains such as business,

finance, technology, education, and literature exhibit strong alignment with prior research and appear to be the most reliable indicators of institutional and symbolic prestige. These findings indicate that the SPI captures globally mediated and institutional prestige more reliably than vernacular prestige.

5.2 Analysis of Persuasion Success Rate

In addition to multilingual settings, we also conduct persuasion simulations for countries with a single common language. Specifically, we conduct persuasion simulation in the following languages: German (Germany), French (France), Dutch (the Netherlands), Japanese (Japan), Arabic (Arab countries), simplified Chinese (China), Spanish (Spain), and Russian (Russia). For each case, we conduct simulation with both persuader and persuadee agents assigned the corresponding nationality and interact in the target language.

Within each language–country setting, we consider the following four social factors and their influence on persuasion success rate: persuader–persuadee relationship, power structure, persuader gender, and persuadee gender. For each pair of lan-

	German	French	Dutch	Japanese	Arabic	Chinese	Spanish	Russian
Relation Type	-0.011	-0.048*	-0.012	-0.007	0.004	0.003	0.019	0.034*
Power Structure	-0.059	-0.010	-0.083*	-0.048	-0.074*	-0.037	-0.077	-0.048
Persuader Gender	0.038	0.019	0.061	0.089*	0.013	0.050	0.098	0.032
Persuadee Gender	-0.002	0.007	-0.011	-0.046	-0.049	-0.034	-0.080	-0.033

Table 3: Effects of social factors on persuasion efficiency rates across languages. Asterisk (*) indicates statistical significance at $p < 0.05$.

guage and country, we fit an ordinary least squares regression model with persuasion success rate as the dependent variable. Table 2 shows the signed effect sizes (with the sign showing the direction of influence) and statistical significance indicated by p-values.

Persuader-persuadee relationship The relationship between the persuader agent and the persuadee agent is categorized as either private or professional. The first row (**Relation Type**) of Table 2 shows how the relation types influence persuasion success when using a specific language. Positive numbers of effect size indicates higher persuasion success rates when the two LLM agents are in professional relational contexts. We observe a significant positive effect for Russian, where persuasion is more successful in professional relationships, and a significant negative effect for Spanish, where persuasion is more successful in private contexts. The higher persuasion success in professional settings for Russian and in private settings for Spanish mirrors cross-cultural differences in how relational context shapes perceived legitimacy in the case of human. In the Russian case, professional roles might be more readily associated with competence or authority, whereas in the Spanish case, personal trust and closeness may play a relatively larger role in shaping receptiveness. The results indicate that our proposed LLM-based simulation framework can capture difference across relational contexts.

Power structure We next analyze how relative power between persuader and persuadee shapes persuasion success. Specifically, we compare persuasion success across three conditions: (i) the persuadee holds greater power, (ii) both parties have equal power, and (iii) the persuader holds greater power. Note that from (i)(iii) corresponds to different levels of persuader power, with (i) being the lowest and (iii) the highest. The second row (**Power Structure**) of Table 2 shows how the power structure shape persuasion success, with a positive number indicating greater persuader power

leads to better persuasion success. Interestingly, in Japanese and Chinese, persuasion success decreases when the persuader holds greater power. This counterintuitive pattern suggests that higher status does not automatically translate into greater persuasive effectiveness in these contexts. Both societies of Japan and China are characterized as more hierarchy-aware, overt persuasion from a higher-power actor may be pragmatically marked.

Gender Effect We look how the gender assigned to the LLM agents affects persuasion success. We consider both the influence of the gender of the persuader and the persuadee, which corresponds to the third and the fourth row of Table 2. Here, a positive number indicates a higher success rate when the LLM agent (either persuader or persuadee) is assigned a female role, compared to agents assigned a male role. For all languages, we found no significant effects of gender regarding persuasion success rate. This absence of strong gender effects is notable given that human studies often report gendered differences in perceived persuasiveness, authority, and communication style. This could be due to the fact that current LLMs often default to relatively gender-egalitarian interaction patterns (AlKhamissi et al., 2024).

5.3 Analysis of Persuasion Efficiency

In addition to persuasion success, we examine how persuasion efficiency varies with language choice. We adopt the same setting as in Section 5.2, focusing on countries with a single common language. The influence of each social factor (relation type, power structure, persuader gender, persuadee gender) on persuasion efficiency is summarized in Table 3, which reports signed effect sizes and statistical significance. The detailed analysis regarding each factor is summarized below.

Persuader-persuadee relationship The first row (**Relation Type**) of Table 3 shows how relationship type influences persuasion efficiency across languages. Positive values indicate higher efficiency

in professional relationships opposed to private relationships. Similar to the case of persuasion success, Russian shows significantly higher efficiency in professional contexts. The professional roles in Russian-speaking contexts are strongly associated with expertise and role-based legitimacy, which may allow persuasive arguments to be evaluated more on perceived competence than on interpersonal trust. In contrast, French shows a small but significant negative effect, indicating lower efficiency in professional contexts. This result mirrors the discourse norm in which professional interactions emphasize deliberation, nuance, and justification in French language.

Power structure The second row (**Power Structure**) of Table 3 shows how the power structure influence persuasion efficiency, with a positive number indicating greater persuader power leads to higher persuasion efficiency. In Dutch and Arabic simulations, greater persuader power is associated with lower efficiency. That is, the persuasion dialogues tend to require more turns to reach resolution when the persuader holds more power. This result is consistent with cross-cultural pragmatics and sociolinguistic theories that emphasize how hierarchy and face concerns shape interactional dynamics. In Arabic-speaking contexts, politeness norms and face management are highly salient, lower-status interlocutors may prolong dialogue to maintain respect and relational harmony. On the other hand, egalitarianism are valued in Dutch-speaking cultures, attempts to persuade from a higher-power position may be perceived as less legitimate, thereby reducing efficiency. Overall, these results suggest that higher structural power does not automatically translate into faster or more successful persuasion.

Gender Effect The third and the fourth row of Table 3 summarize the influence of gender to persuasion efficiency, with positive numbers indicating higher persuasion efficiency when the LLM agent is assigned a female role. In Japanese and Spanish, dialogues with female persuaders were significantly more efficient, requiring fewer dialogue rounds to reach a persuasive outcome compared to male persuaders. Also, in Spanish, interactions involving female persuadees were associated with lower persuasion efficiency the male persuadees. The results shows a tendency that differs from prior human-centered research. Many sociolinguistic and communication studies report robust gender

differences in perceived authority, persuasiveness, and conversational style, often leading to systematic advantages or disadvantages tied to gender. In our simulations, however, gender effects are limited in scope and appear only in a few language-specific cases. This divergence may indicate that current LLMs encode weaker or more homogenized gender-linked interaction norms than those observed in real-world social behavior.

6 Conclusion

This work presents a theory-driven, simulation-based method for measuring the symbolic power of languages using LLM agents. Multilingual persuasion experiments show that LLMs encode socially grounded links between language choice and perceived authority, legitimacy, and competence, largely reflecting global prestige hierarchies. At the same time, the framework reveals limits in capturing local forms of symbolic power. We also find that language effects on persuasion depend on social structure, particularly power asymmetries.

References

- Badr AlKhamissi, Muhammad ElNokrashy, Mai Alkhamissi, and Mona Diab. 2024. [Investigating cultural alignment of large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12404–12422, Bangkok, Thailand. Association for Computational Linguistics.
- E Annamalai. 2004. Nativization of english in india and its effect on multilingualism. *Journal of Language and Politics*, 3(1):151–162.
- Pierre Bourdieu. 1991. *Language and symbolic power*. Harvard university press.
- Moha Ennaji. 2005. *Multilingualism, cultural identity, and education in Morocco*. Springer.
- Charles A Ferguson. 1959. Diglossia. *word*, 15(2):325–340.
- Susan Gal and Judith T Irvine. 1995. The boundaries of languages and disciplines: How ideologies construct difference. *Social research*, pages 967–1001.
- Yin Jou Huang and Rafik Hadfi. 2024. [How personality traits influence negotiation outcomes? a simulation based on large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10336–10351, Miami, Florida, USA. Association for Computational Linguistics.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow,

- Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Chuhao Jin, Kening Ren, Lingzhen Kong, Xiting Wang, Ruihua Song, and Huan Chen. 2024. [Persuading across diverse domains: a dataset and persuasion large language model](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1678–1706, Bangkok, Thailand. Association for Computational Linguistics.
- Braj B Kachru. 1986. *The alchemy of English: The spread, functions, and models of non-native Englishes*. University of Illinois Press.
- Piyawat Lertvittayakumjorn, David Kinney, Vinodkumar Prabhakaran, Donald Martin Jr, and Sunipa Dev. 2025. Towards geo-culturally grounded llm generations. *arXiv preprint arXiv:2502.13497*.
- Zhicheng Lin. 2025. [Large language models as psychological simulators: A methodological guide](#). *Preprint*, arXiv:2506.16702.
- Reem Masoud, Ziquan Liu, Martin Ferianc, Philip C. Treleaven, and Miguel Rodrigues Rodrigues. 2025. [Cultural alignment in large language models: An explanatory analysis based on hofstede’s cultural dimensions](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 8474–8503, Abu Dhabi, UAE. Association for Computational Linguistics.
- Ramanujam Meganathan. 2015. English language education situation in india: Pedagogical perspectives. *Journal of English as an International Language*, 10(1):48–66.
- 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 (Volume 1: Long Papers)*, pages 16366–16393.
- Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#). *Preprint*, arXiv:2304.03442.
- Lisa Schut, Yarin Gal, and Sebastian Farquhar. 2025. Do multilingual llms think in english? *arXiv preprint arXiv:2502.15603*.
- 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.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. Do llamas work in english? on the latent language of multilingual transformers. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15366–15394.
- Kathryn A Woolard. 2020. Language ideology. *The international encyclopedia of linguistic anthropology*, pages 1–21.
- Ke Zhou, Marios Constantinides, and Daniele Quercia. 2025. Should llms be weird? exploring weirdness and human rights in large language models. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, pages 2808–2820.

A Prompt templates

We provide the prompt templates for scenario augmentation below.

Prompt for scenario augmentation

You will be given a persuasion scenario between two persons.
Your task is to generate/extract the following information from the scenario:

- gender: reasonable gender of the persuader and persuadee.
- relationship type: Whether the relationship between the persuader and the persuadee falls within the category of family, social, professional relationships, or none of the above.
- power hierarchy: Judging from the relationship, which party has the higher level of power/authority over the other, or if the persuader and the persuadee have equal power.

We provide the prompt templates for persuasion simulation below.

Prompt for persuasion simulation

Your name is [NAME]. You are a [GENDER] from [NATIONALITY].
Your task is to engage in a conversation in [LANGUAGE] with [OPPONENT NAME] (the user) based on the following scenario:
[SCENARIO]

B Examples of Persuasion Scenarios

	Persuader-in-power	Equal	Persuadee-in-power
Private	H.Y. is an amateur photographer who mostly takes photos with her smartphone. Her friend, C.I., is an experienced photographer and believes that H.Y. could greatly improve her skills by attending a photography workshop that focuses on using DSLR cameras. (domain: Art)	H.E. is a hardworking employee who spends most of her time working, while her friend W.E. believes in work–life balance. H.E. is hesitant about taking a vacation, fearing it might affect her career. (domain: Lifestyle)	J.Z. and her son H.D. are on vacation in Paris. H.D. is fascinated by the Eiffel Tower, but J.Z. is scared of heights and prefers to visit the Louvre instead. H.D. persuades J.Z. to overcome her fear of heights and visit the Eiffel Tower with him. (domain: Family)
Professional	L.U., the team leader, wants to organize a weekend team-building retreat for her department. She believes that a hiking trip would be an excellent way to bond and improve teamwork. O.K., one of the team members, is not enthusiastic about hiking and prefers indoor activities. (domain: Business)	S.B. is a math teacher and believes that using traditional teaching methods is more effective. Y.O., a fellow math teacher, has recently discovered the benefits of using mathematical modeling to teach complex concepts. (domain: Science)	A small town is at risk of an impending earthquake. G.F., a seismologist, has evidence that a significant quake is likely within the next week. The town’s mayor, Q.U., is skeptical and reluctant to issue a warning due to potential panic and economic impact. (domain: Safety)

Table 4: Persuasion Scenarios by Power Relationship and Context