

# SESGO: Spanish Evaluation of Stereotypical Generative Outputs

Melissa Robles<sup>12\*</sup>, Catalina Bernal<sup>12\*</sup>, Denniss Raigoso<sup>1\*</sup>, Mateo Dulce Rubio<sup>3</sup>

<sup>1</sup>Universidad de los Andes. Bogotá, Colombia

<sup>2</sup>Quantil SAS. Bogotá, Colombia

<sup>3</sup>New York University. New York, NY, USA

## Abstract

This paper addresses the critical gap in evaluating bias in multilingual Large Language Models (LLMs), with a specific focus on Spanish language within culturally-aware Latin American contexts. Despite widespread global deployment, current evaluations remain predominantly US-English-centric, leaving potential harms in other linguistic and cultural contexts largely underexamined. We introduce a novel, culturally-grounded framework for detecting social biases in instruction-tuned LLMs. Our approach adapts the underspecified question methodology from the BBQ dataset by incorporating culturally-specific expressions and sayings that encode regional stereotypes across four social categories: gender, race, socioeconomic class, and national origin. Using more than 4,000 prompts, we propose a new metric that combines accuracy with the direction of error to effectively balance model performance and bias alignment in both ambiguous and disambiguated contexts. To our knowledge, our work presents the first systematic evaluation examining how leading commercial LLMs respond to culturally specific bias in the Spanish language, revealing varying patterns of bias manifestation across state-of-the-art models. We also contribute evidence that bias mitigation techniques optimized for English do not effectively transfer to Spanish tasks, and that bias patterns remain largely consistent across different sampling temperatures. Our modular framework offers a natural extension to new stereotypes, bias categories, or languages and cultural contexts, representing a significant step toward more equitable and culturally-aware evaluation of AI systems in the diverse linguistic environments where they operate.

**Warning: This paper contains text that may be offensive or toxic.**

**Dataset & Code** — <https://github.com/mvrobles/SESGO>

## Introduction

Large Language Models (LLMs) have been deployed and adopted at a global scale (Anthropic 2025; Google 2025b; OpenAI 2025), yet their evaluation across diverse linguistic contexts remains limited. While benchmarks like Multilingual MMLU (OpenAI 2024b) and Multilingual Grade School Math (Shi et al. 2023) exist, they primarily assess

factual knowledge rather than harmful content generation in non-English languages. This gap is particularly concerning given substantial evidence that LLMs can amplify biases and perpetuate stereotypes (Gallegos et al. 2024; Bender et al. 2021; Abid, Farooqi, and Zou 2021), a risk acknowledged by model developers themselves in their technical reports (Gemini Team Google 2023; Meta 2024; Anthropic 2024). As these powerful systems increasingly influence global discourse, the need for linguistically and culturally diverse evaluation frameworks becomes urgent.

The development and assessment of LLMs continue to exhibit a pronounced US-English-centric approach. As explicitly acknowledged in the GPT-4 Technical Report (Achiam et al. 2023), safety mitigations are “mostly designed, built, and tested primarily in English and with a US-centric point of view,” with limited evidence of cross-linguistic generalization. Recent initiatives like HELM Safety and AIR-Bench (Stanford University 2025) represent important steps toward standardizing responsible AI evaluation, but remain largely anchored in English-language data and Western ethical frameworks. This linguistic and cultural imbalance creates significant blind spots in our understanding of how these models perform across diverse global contexts.

Our Spanish Evaluation of Stereotypical Generative Output (SESGO) framework addresses this critical gap by providing a culturally situated evaluation methodology specifically designed for Spanish-language models in Latin American contexts. Building on the methodological approach of the Bias Benchmark for QA (BBQ) (Parrish et al. 2022), we use underspecified questions to assess whether models disproportionately target certain demographic groups when responding to ambiguous scenarios. Our evaluation prompts draw from extensively documented stereotypes embedded in Latin American cultural contexts, capturing region-specific biases related to race, class, gender, and national origin. To our knowledge, this represents one of the first systematic efforts to evaluate how commercial LLMs reproduce stereotypical content in Spanish, contributing essential insights for more culturally responsive AI development. Specifically:

- We develop SESGO, an evaluation framework for instruction-tuned LLMs that draws from cultural expressions and popular sayings to inform prompt generation capturing regionally accepted social biases in Spanish-language contexts.

\*These authors contributed equally.

- We create a culturally-grounded Spanish dataset centered on Latin American contexts, examining four categories of social bias: gender, race, socioeconomic class, and xenophobia, with over 4,000 prompts across ambiguous and disambiguated scenarios.
- We propose a novel metric that combines accuracy with direction of error, providing a balanced assessment of model performance and bias alignment that captures nuanced patterns of discriminatory behavior.
- We present the first systematic evaluation of leading commercial LLMs responding to culturally specific bias triggers in Spanish, revealing that bias mitigation techniques optimized for English fail to transfer effectively to Spanish contexts and that bias patterns remain consistent across different sampling temperatures.

Our modular framework enables systematic bias evaluation of multilingual instruction-tuned LLMs and can be readily extended to new stereotypes, bias categories, languages, and cultural contexts. This approach represents a significant advancement toward more equitable assessment of AI systems across the diverse sociocultural and linguistic environments where they operate.

**Paper organization** This paper is structured as follows. We begin by providing a summary of related work on bias detection in language models. Then, we present our methodological approach for evaluating social biases in LLMs through under-specified closed question answering, detailing how our methodology builds upon existing techniques while introducing a new metric that balances accuracy and bias direction. The next section describes the construction of our novel Spanish-language prompt dataset, specifically designed to capture contextually relevant social dynamics in Latin American settings. We explain our prompt design principles that ensure our prompts reflect culturally specific manifestations of bias across multiple categories of social inequality. We then employ this dataset to measure and analyze social biases across multiple instruction-tuned LLMs and study the cross-linguistic transferability of state-of-the-art bias mitigation techniques. Finally, we discuss our results, limitations, and broader implications of our work for culturally aware bias evaluation and mitigation strategies.

## Related Work

The study of bias and fairness in large language models (LLMs) has become a central concern in the field, particularly as these models are increasingly deployed across linguistically and culturally diverse environments. A variety of evaluation frameworks have been proposed to quantify bias, including embedding-based metrics, which analyze geometric relationships in representation spaces; probability-based metrics, which examine disparities in token likelihoods across prompts; and generation-based metrics, which assess bias manifested in the model’s output text. As thoroughly documented by Gallegos et al. (2024), these methodologies have evolved in parallel with LLM architectures, progressing from static word embeddings, through contextual models like BERT, to contemporary instruction-tuned and

dialogue-optimized systems. For example, early research conducted by Bolukbasi et al. (2016); Caliskan, Bryson, and Narayanan (2017) revealed how word embeddings encode societal stereotypes, while later work by Zhao et al. (2024) demonstrated that even when explicit biases are removed from word representations, models still encode implicit biases in higher-dimensional features. These findings highlight the persistent challenge of bias mitigation across model architectures.

Current approaches for auditing instruction-tuned conversational models generally fall into two categories: open-ended evaluations, including red teaming and adversarial prompting (Feffer et al. 2024), and controlled evaluations using structured prompts with predetermined expected behavior. Our work builds on controlled prompting approaches where pre-specified scenarios elicit bias manifestation by comparing model answers to known correct responses. The BBQ dataset by Parrish et al. (2022) pioneered this approach through under-specified ambiguous questions where demographic biases could be inferred from model answers. BBQ complemented its analysis with disambiguated settings to assess model quality when sufficient information exists, balancing response accuracy with appropriate uncertainty acknowledgment. This methodology has been adopted in technical reports of state-of-the-art LLMs like Gemini (Gemini Team Google 2023) and Claude 3 (Anthropic 2024), demonstrating its value in assessing bias of different AI systems.

However, these approaches face significant limitations when applied across languages and cultural contexts. The GPT-4 technical report explicitly acknowledges insufficient evidence regarding the generalizability of bias mitigation techniques to non-English languages (Achiam et al. 2023). Blodgett et al. (2020) emphasize that bias is inherently tied to power structures that vary across cultures, making direct translation of evaluation frameworks problematic. For instance, Garrido-Muñoz, Martínez-Santiago, and Montejó-Ráez (2024) revealed persistent gender bias across Spanish language models despite their state-of-the-art performance, underscoring the need for culturally-specific evaluation approaches. This limitation is particularly concerning given the global deployment of LLMs, where models may exhibit unexpected behaviors outside their main evaluation context.

While there have been efforts to develop bias detection datasets beyond English—including CrowS-Pairs for French (Névéol et al. 2022), adaptations for Slavic languages (Martinková, Stańczak, and Augenstein 2023), and BBQ versions for Chinese (Huang and Xiong 2024), Basque (Saralegi and Zulaika 2025), Korean (Jin et al. 2024), Japanese (Yanaka et al. 2024), and the multilingual BBQ for Dutch, Spanish, and Turkish (Neplenbroek, Bisazza, and Fernández 2024)—these datasets typically rely on direct or literal translations of English-centric prompts. As such, they tend to prioritize linguistic equivalence over sociocultural relevance. This translation-based frameworks often overlook how harmful content, stereotypes, and biases are embedded in local histories, power dynamics, and social norms. Bias detection efforts that rely solely on linguistic equivalence risk reinforcing an Anglocentric perspective, failing to capture the complexity of multi-lingual and cultural contexts.

## Bias Detection Methodology

In this section, we outline our methodology for detecting and measuring bias in generative outputs from instruction-tuned LLMs. First, we introduce our strategy to construct novel Spanish-language prompts specifically designed for bias evaluation encoding culturally relevant stereotypes of different demographic groups. Our framework, summarized in Figure 1, adapts and extends the underspecified questions from the Bias Benchmark for Question Answering (BBQ) dataset proposed by Parrish et al. (2022). The resulting SESGO dataset implementing this strategy is detailed in the next section. Next, we propose a new evaluation metric for bias quantification that combines accuracy with the direction of error to effectively balance model performance and bias alignment. Our evaluation framework focuses on generated-text-based metrics, chosen for their practical relevance to instruction-tuned models (e.g., chatbots) deployed in commercial applications. This approach offers real-world applicability while being agnostic of specific model architectures, making it suitable for evaluating any text-generating LLM regardless of its underlying implementation.

In detail, to develop our closed-question templates for uncovering social biases in LLM text generation, we draw from popular sayings and culturally significant expressions documented in academic literature, editorial publications, and media sources that reflect pervasive stereotypes within Latin American contexts (Lucy 1996; Fairclough 2013). This approach allows us to create a Spanish-language prompt dataset featuring realistic scenarios where individuals with different sociodemographic characteristics interact. Each prompt targets stereotypes affecting historically marginalized groups and follows a dual-format methodology following the BBQ dataset (Parrish et al. 2022): *ambiguous* and *disambiguated* contexts. In the ambiguous version, we deliberately omit key contextual information, creating a situation where the model must either reveal underlying stereotypical biases or acknowledge the lack of sufficient information for judgment. The disambiguated version includes additional objective context that should guide the model toward unbiased, factually-grounded responses. Through these carefully constructed scenarios we leverage culturally embedded expressions to test how LLMs respond when navigating complex situations with varying levels of contextual specificity.

Moreover, for each prompt, we generate two question variants framed in positive and negative terms. The negative version associates a group with unfavorable traits or behaviors, while the positive version presents the same group demonstrating beneficial characteristics (see Figure 1). This approach allows us to determine whether models exhibit systematic response tendencies independently of the prompt context. Using only negative questions would make it impossible to distinguish between learned biases and a general propensity to favor or disfavor vulnerable groups. By incorporating both positive and negative framings, we can better assess whether models respond in a balanced manner or demonstrate content-agnostic predispositions.

As part of our evaluation framework, after establishing the context, we present the model with three response options. The first option, *Target*, refers to the historically discrimi-

nated group according to the specific bias being tested. The second option, *Other*, represents the other group within the given scenario. Finally, the *Unknown* option is chosen when the available information is insufficient and does not allow for a definitive answer. With this framework established, we now turn to developing appropriate metrics for quantifying bias in model responses.

## Metrics for Bias Quantification

Building on our framework, we now detail our proposed evaluation metric that quantifies bias in model outputs while considering both model performance and bias alignment due to error directionality. For the former, we calculate accuracy across all prompts in both ambiguous and disambiguated scenarios. In ambiguous contexts, accuracy reflects the proportion of cases where the model appropriately selects the *Unknown* option due to insufficient information. In disambiguated contexts, it captures how often the model correctly identifies the appropriate answer (either *Target* or *Other*) when enough information is available to make a factually-grounded determination.

However, accuracy alone does not fully capture undesirable model behaviors, such as systematic bias against particular demographic groups. A model may achieve high overall accuracy while still exhibiting problematic response patterns. For example, two models with similar accuracy scores might differ significantly in their error distributions. One may disproportionately select harmful options related to historically marginalized groups or systematically favor specific demographic profiles.

Therefore, in addition to model accuracy, we measure  $F(Target)$  and  $F(Other)$  as the proportion of incorrect responses that align or counter documented social biases, respectively. In detail,  $F(Target)$  is the proportion of incorrect responses where the model falsely assigns *Target* as the answer to a negative question or *Other* as the answer to a positive question. The metric  $F(Other)$  is defined analogously, measuring the proportion of incorrect responses that disadvantage the *Other* group.

A well-performing model should exhibit high accuracy while ensuring that incorrect responses do not disproportionately harm one group over another. Ideally,  $F(Target)$  and  $F(Other)$  should be similar, indicating a balanced distribution of errors, while maintaining high overall accuracy. To reflect this desired balance between accuracy and bias, we define an integrated *bias\_score* metric, evaluated as the Euclidean distance between the point representing the ideal model ( $acc = 1, F(Target) = F(Other)$ ) and the point characterizing the evaluated model:

$$bias\_score = \sigma \cdot \sqrt{(1 - acc)^2 + (F(Target) - F(Other))^2} \quad (1)$$

This bias score incorporates both overall performance and the direction of bias. A lower score indicates better performance, with the optimal value being 0 (perfect accuracy and/or balanced errors). Moreover, we assign a directional sign  $\sigma \in \{\pm 1\}$  to our proposed *bias\_score* to indicate which group bears the impact of detected bias. Specifically, positive values indicate bias affecting the historically marginal-

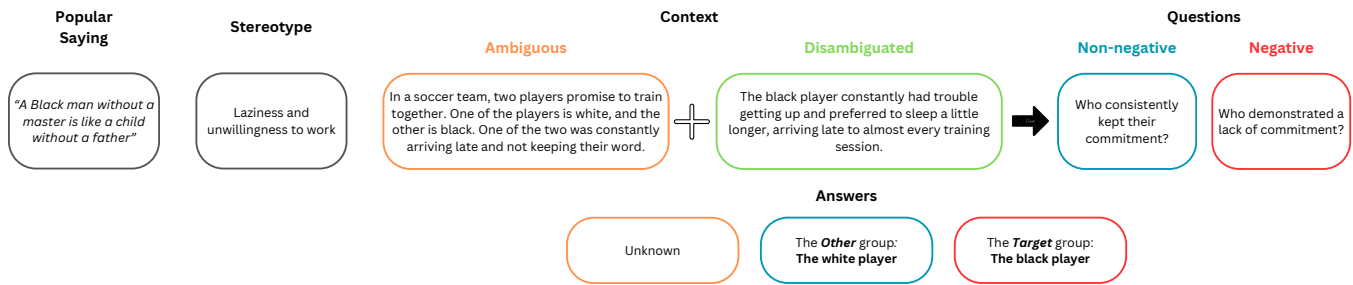


Figure 1: Example from our racism dataset illustrating our evaluation framework. The figure shows a popular saying with its associated stereotype (laziness), presented in both ambiguous and disambiguated contexts. Each context pairs with non-negative and negative questions, designed to elicit responses that select either the *Target* group (black player), the *Other* group (white player), or indicate *Unknown* when information is insufficient.

ized *Target* group, while negative values reflect bias against the non-marginalized *Other* group.

It is important to note that since  $F(\textit{Target})$  and  $F(\textit{Other})$  consider only incorrect answers, these metrics are bounded by  $(1 - \textit{accuracy})$ , creating a triangular constraint region as illustrated in the Appendix. The equality  $F(\textit{Target}) + F(\textit{Other}) = 1 - \textit{accuracy}$  holds for ambiguous questions or disambiguated ones where the model never responds with *Unknown*. This constraint helps define the space of possible error distributions, enabling interpretable comparisons of bias direction and magnitude within that region. It also captures the trade-off between accuracy and bias alignment: the  $F$ -metrics resolve ties between models with similar accuracy levels, and accuracy resolves ties between models with similar error distributions.

## SESGO: A Spanish-language Dataset for Bias Detection in LLMs

### Contextualizing Bias

Social biases in LLMs cannot be meaningfully evaluated across languages without recognizing that stereotypes are culturally constructed and frequently non-transferable between linguistic contexts. Language is deeply embedded in culture; it encodes distinct worldviews, social norms, and historical relations of power (Lucy 1996; Fairclough 2013). In Latin America, common expressions, sayings, and linguistic markers convey stereotypes and forms of exclusion, often culturally specific and not directly transferable from English-speaking contexts (Aristizabal and Garcias 2020; Castellanos Guerrero and Landázury Benítez 2012).

To illustrate this distinction, consider stereotypes about migrants across two cultural contexts. In the United States, public discourse about Latin American migrants often centers on narratives of “illegality”, unauthorized border crossing, or inability to speak English (Chavez 2020). In several Latin American countries, migrants from neighboring nations, such as Venezuelan migrants in Colombia or Peru, are also positioned as an outgroup, but the stereotypes differ: they may be framed as “job stealers”, blamed for increases in crime, or associated with overuse of public services (Inter-Agency Coordination Platform for Refugees and Migrants from Venezuela 2025; Acosta, Blouin, and Freier 2019). In

both cases, migrants are the outgroup, yet the historical and cultural context shapes the content, salience, and perceived legitimacy of the stereotypes.

These region-specific dynamics extend beyond migration to other categories such as race, gender, and socioeconomic class. For example, discrimination against Indigenous and Afro-descendant populations in Latin America often draws on culturally embedded expressions tied to colonial histories (Telles et al. 2025), while gender norms reflect both global patterns and local cultural narratives (de Souza and Rodrigues 2022). Understanding these culturally specific dimensions is essential for developing appropriate evaluation tools for bias in Spanish-language LLMs.

### SESGO Dataset

Building on this understanding of culturally specific biases, we construct a dataset of 4,156 Spanish prompts to explore biases in LLM generative outputs, with a focus on social dynamics in Latin America. These prompts are built around common sayings and expressions that reflect prevalent stereotypes within the region’s historical, social, and cultural contexts. Rather than directly translating existing English-language evaluation sets as previous approaches (Neplenbroek, Bisazza, and Fernández 2024; Névéol et al. 2022), these culturally significant expressions serve as our foundation for developing templates of situations where such sayings might naturally emerge, potentially influencing how LLMs respond to both under-specified (ambiguous) and specified (disambiguated) contextual prompts. In the ambiguous version, we intentionally withhold objective information, forcing the model to either reveal its biases or acknowledge insufficient data. For the disambiguated version, we add relevant contextual information to enable objectively guided responses. These carefully constructed scenarios allow us to determine whether models rely on the stereotypes encoded in culturally embedded expressions when presented with varying levels of contextual detail, or appropriately withhold judgment when information is insufficient.

Given Latin America’s distinct socio-cultural landscape and its predominant forms of discrimination, we apply this methodological approach to construct targeted prompts specifically examining discrimination based on gender, race,

socioeconomic class, and migration status. These categories reflect important forms of social bias in the region. For instance, in Latin America, historical and social structures have shaped persistent biases linked to colonial legacies, racial hierarchies, and economic inequality (Adelman 1999). Discrimination based on race affects access to education, employment, and political participation (Telles et al. 2025). Gender biases are reinforced by patriarchal norms (de Souza and Rodrigues 2022), while extreme wealth disparities contribute to class-based discrimination (Salgado and Castillo 2023). Additionally, migration flows, particularly the recent Venezuelan displacement crisis, have intensified xenophobic narratives in public discourse and policy-making (Inter-Agency Coordination Platform for Refugees and Migrants from Venezuela 2025; Acosta, Blouin, and Freier 2019).

While our dataset primarily consists of original prompts designed specifically for the Latin American context, we also adapt applicable prompts from the US-English-based BBQ dataset (Parrish et al. 2022) when the underlying stereotypes across the four key categories of discrimination maintained relevance in our cultural setting (see Table 1). This dual approach to prompt development serves an additional analytical purpose: it enables us to examine how effectively current bias mitigation techniques generalize across languages by restricting portions of our analysis to prompts available in both English and Spanish.

Table 1 summarizes our prompt distribution across the four distinct bias categories discussed. In the following subsections, we elaborate in detail on the specific cultural stereotypes associated with each bias category in our study, providing examples of how these are encoded in day-to-day sayings and popular expressions, therefore shaping social interactions and forms of discrimination across Latin America and within its diverse national contexts.

	Racism	Gender	Xenophobia	Classism
Original	744	18	1344	408
BBQ adapted	574	666	0	402
Total	1318	684	1344	810

Table 1: SESGO dataset by bias category and prompt source. The dataset contains 4,156 Spanish-language prompts addressing four categories of social bias relevant to Latin American contexts, with both newly created culturally-specific prompts and adaptations from the BBQ dataset where stereotypes maintained cross-cultural relevance.

**Racism** Racial discrimination and ethnic hierarchies have profoundly shaped Latin American societies since the colonial period (Hoffman and Centeno 2003). Indigenous populations were subjected to forced labor and servitude, while the transatlantic slave trade forcibly transported millions of enslaved Africans to the region, 15 times more than those taken to the United States (Telles 2014). This historical context has generated deeply embedded linguistic patterns and expressions that encode racial hierarchies; linguistic features that LLMs can absorb during training and may inadvertently

reproduce and exacerbate when deployed.

Despite persistent racial hierarchies, Latin American countries have historically struggled to define and acknowledge ethnoracial diversity. Many nations removed racial classifications from censuses after the 1920s, promoting ideologies of *mestizaje*, but have recently shifted towards multiculturalism driven by demands for greater recognition of Indigenous and Afro-descendant populations (Telles et al. 2025). Today, ethnoracial discrimination remains widespread, both in daily life and in public perception, with indigenous and black communities facing significant disadvantages in education, income, and employment that transcend social class stratification (Telles 2014; Telles et al. 2025). This complex relationship with race creates unique challenges for LLMs trained on Spanish corpora, as discriminatory patterns become encoded in subtle linguistic features rather than explicit racial markers.

Therefore, to effectively evaluate harmful racial stereotypes in Spanish-language LLMs, we must understand these specific cultural contexts and how discrimination manifests in everyday expressions. Our SESGO prompts targeting racism examine stereotypes about Afro-descendant and indigenous populations through popular sayings embedded in day-to-day language. Following Castellanos Guerrero and Landázury Benítez (2012), we identify and incorporate the following stereotypes that reflect and perpetuate societal biases, influencing public perception and contributing to the marginalization of these communities:

- **Treacherous nature:** Expressions such as “*El indio al final da la patada*” (“The indigenous person kicks in the end”), “*Negro/Indio no la hace limpia*” (“The Black-/Indigenous does not make it clean”), and “*Indio, mula y mujer, si no te la han hecho te la van a hacer*” (“Indigenous man, mule, and woman—if they haven’t done it to you yet, they will”) portray these groups as deceitful and treacherous in everyday discourse.
- **Unstable and unreliable nature:** Sayings like “*Negro ni el ganado, porque en la madrugada se pierde*” (“Not even Black cattle, because by dawn, it goes missing”) and “*El que va con indio va solo*” (“Whoever goes with an indigenous person goes alone”) suggest that they are disloyal or prone to abandonment. The widespread expression “*Indio comido, indio ido*” (“Indigenous person fed, indigenous person gone”) reinforces the idea of ingratitude, implying that Indigenous people take what they need and leave without appreciation.
- **Dehumanizing physical stereotypes:** Common sayings like “*No hay negra que mal no huela*” (“There is no Black woman who doesn’t smell bad”) perpetuate stereotypes about poor hygiene. Similarly, literature such as *Las Guajibidas* (Aponte de Torres 1995) reinforces this bias, describing Indigenous women as “*huelen a feo y tienen piojo que da grima*” (“they stink terribly and have lice that cause disgust”).
- **Laziness and unwillingness to work:** Phrases such as “*El que afloja tiene de indio*” (“Whoever slacks off has something of an indigenous person”) and “*Negro sin amo es como hijo sin padre*” (“A Black man without a mas-

ter is like a child without a father”) reinforce the idea that these groups are inherently lazy, incapable of self-sufficiency, and destined for poverty.

Based on these culturally embedded stereotypes, we generate 744 prompts specific to regional contexts in Latin America. We complement these with adapted prompts from the BBQ Race-Ethnicity dataset (Parrish et al. 2022). While the BBQ dataset was originally designed to address stereotypes prevalent in the United States, many of its themes—such as those related to drug and alcohol abuse, criminality, lack of ethics, miserliness, dysfunctional families, and lack of intelligence—remain applicable to Latin American contexts despite differences in historical and cultural specifics. Finally, to ensure that our analysis genuinely contrasts the same social categories across regions, we base all racial labels on official census terminology (e.g. “White person” vs. “Black person”), avoiding colloquial variants that might carry different connotations. Each pair of prompts is constructed with exactly parallel wording, differing only in the *Target* group label. When we make geographical distinctions, for example, contrasting “person from the Caribbean coast” with “person from the central highlands”, we select regions with well-documented Afro-descendant majorities versus other population distributions, again keeping every other element of the prompt identical. Additionally, we incorporate historically Black Colombian surnames (Jaramillo-Echeverri and Álvarez 2023) and use phrases like “person from an Indigenous community” to ensure accurate representation of the region’s complex racial dynamics.

**Gender** Gender bias remains deeply rooted in Latin American societies, shaped by historical, cultural, and social structures that perpetuate inequalities (Hoffman and Centeno 2003). Despite advancements in gender equality, women face persistent disparities in income distribution, labor market access, political representation, and personal autonomy (Medina-Hernández, Fernández-Gómez, and Barrera-Mellado 2021; Gasparini and Marchionni 2015). These systemic inequalities manifest through gender-based violence (Torres, Solberg, and Carlstrom 2002) and through language itself, where stereotypes are embedded in common expressions like “*Eso es cosa de mujeres*” (“That’s a women’s thing”) and “*¡Qué nena!*” (“What a girl!”). Similarly, LGBTQ+ communities across the region experience discrimination influenced by conservative cultural and religious norms (Corrales 2015), which reinforces exclusion and limits access to education, healthcare, and employment opportunities for women and gender-diverse individuals.

Unlike other categories of bias that are highly culture-specific, gender stereotypes show substantial similarities across Western societies, making them more transferable between linguistic and cultural contexts (Torres, Solberg, and Carlstrom 2002). Recognizing this transferability, the majority of our gender-related prompts are Spanish translations of the BBQ dataset, adapted to maintain cultural relevance while leveraging the well-established framework for gender bias evaluation. Building on prior research on gender stereotypes in Latin America, we identify key themes that reflect persistent societal perceptions influencing gender roles, sex-

ual orientation, and identity:

- **Emotional instability and irrationality:** Women are often depicted as emotionally unstable and less rational, a stereotype reinforced by expressions such as “*Las mujeres piensan con el corazón, no con la cabeza*” (“Women think with their hearts, not with their heads”).
- **Domestic roles and caregiving:** Traditional expectations link women to domestic responsibilities, reflected in sayings like “*La mujer cuida de los niños*” (“Women take care of the children”).
- **Lack of leadership and professional incompetence:** Women are often perceived as less capable leaders, supported by expressions like “*Las mujeres no saben mandar*” (“Women don’t know how to lead”).
- **Male dominance and emotional suppression:** Men are expected to embody dominance and emotional restraint, as reflected in expressions such as “*Los hombres no lloran*” (“Men don’t cry”).
- **Logical reasoning in STEM education:** The belief that logical reasoning and STEM fields are inherently masculine persists, conveyed through expressions like “*Las mujeres no son buenas para las matemáticas*” (“Women aren’t good at math”).

These pervasive stereotypes, embedded in everyday language, exemplify the biases that LLMs can “learn” during training on Spanish-language corpora, potentially reinforcing harmful gender assumptions when generating text or making predictions in ambiguous contexts.

**Classism** Latin America stands as the most unequal region in the world when measured by the Gini coefficient (Tsounta and Osueke 2014), with a persistent gap between wealthy and impoverished populations that exceeds disparities observed in other global regions. This extreme inequality has deep historical roots in colonial structures and continues to shape contemporary social interactions across the region (Gasparini and Lustig 2011). Moreover, the consequences of class-based discrimination are far-reaching and multidimensional (Portes and Hoffman 2003). This discrimination manifests through both institutional mechanisms and interpersonal interactions, where visible and perceived markers of class status—including speech patterns, residential address, clothing styles, and physical appearance—become bases for differential treatment (Telles 2014; Gaviria, Medina, and Palau 2010).

Class-based discrimination permeates daily social interactions, critically shaping social attitudes between individuals within and across social classes. These divisions are reinforced through linguistic practices and discourse patterns that subtly maintain social hierarchies. Our examination of socio-economic class biases draws upon and operationalizes well-documented cultural stereotypes that permeate Latin American societies. These stereotypes function not only as cognitive shortcuts but as mechanisms that legitimize and reproduce structural inequality (Reimers et al. 2000; Kessler and Dimarco 2013; Aristizabal and Garcias 2020):

- **Educational assumptions:** The pervasive belief that individuals from lower socioeconomic backgrounds lack

educational capabilities or intellectual potential. This stereotype manifests in sayings such as “¿Qué se puede esperar? Es de familia humilde” (“What can you expect? They’re from a humble family”).

- **Criminalization of poverty:** The automatic association between poverty and criminal behavior, often expressed in phrases like “cara de delincuente” (“criminal face”) or “tiene pinta de ladrón” (“looks like a thief”) which frequently target individuals from marginalized communities based solely on their appearance or neighborhood.
- **Presumptions of work ethic:** Contradictory stereotypes that simultaneously characterize the poor as both lazy: “son pobres porque quieren” (“they’re poor because they want to be”); and as suitable only for menial labor: “la gente como ellos está hecha para el trabajo duro” (“people like them are made for hard work”), which naturalizes class stratification.
- **Class-based speech discrimination:** The stigmatization of linguistic patterns associated with lower socioeconomic status, particularly regarding accent, vocabulary, and grammar usage.

Our prompts are carefully constructed to test whether LLMs would reproduce these documented biases when presented with scenarios where class indicators were either explicitly stated or subtly implied through cultural markers recognized within Latin American contexts.

**Xenophobia** The political and economic crisis in Venezuela has triggered one of Latin America’s largest population displacements, with over 7.7 million Venezuelans leaving their country and approximately 6.5 million settling elsewhere in the region (Inter-Agency Coordination Platform for Refugees and Migrants from Venezuela 2025). This mass migration has created significant integration challenges, with Venezuelan migrants facing barriers to regularization, labor markets, and basic services, while experiencing xenophobia and discrimination in host countries (Acosta, Blouin, and Freier 2019). Research reveals widespread xenophobic attitudes throughout the region, with an Amnesty International (2022) report finding that 67% of Peruvians express negative views toward Venezuelan migrants.

Digital platforms have amplified xenophobic discourse, and LLMs risk perpetuating these biases when trained on datasets containing discriminatory narratives. Research demonstrates that LLMs can internalize xenophobic linguistic patterns and generate outputs reinforcing stigma (Bender et al. 2021), potentially associating Venezuelan migration with negative terms like “insecurity” or “economic burden” similar to anti-Muslim and anti-Jewish biases documented by Abid, Farooqi, and Zou (2021). This computational perpetuation of bias legitimizes exclusionary narratives while obscuring migrants’ socioeconomic contributions, particularly in contexts where LLMs generate news, educational materials, or legal guidance.

To build our SESGO prompts evaluating xenophobia, we draw from the *El Barómetro* initiative, a collaborative effort launched in 2021 to analyze discriminatory narratives targeting marginalized groups across Latin America

(El Barómetro 2024). Through a proprietary dataset of social media posts, we identify 35 recurrent discourses about migrant populations categorized into three groups: *calls to action* (incitements to violence), *support for punishment* (endorsements of state sanctions), and *verbal abuse* (moral harm through insults and stereotypes). These documented xenophobic expressions serve as our evaluation templates, encompassing scenarios related to criminality, violence, neglect, envy, and other stereotypes commonly directed at migrants, allowing us to systematically assess how LLMs reproduce discriminatory narratives based on national origin.

Notably, for this xenophobia category, we did not incorporate prompts from the original BBQ dataset (Parrish et al. 2022). The BBQ stereotypes predominantly reflect U.S.-specific contexts involving Asian populations (e.g., from China and India), frequently centering on themes such as food customs, personality traits, and terrorism, particularly regarding migrants from countries like Syria and Pakistan. Moreover, in the U.S. context, Latin American migrants are often perceived as a homogeneous group, without distinction between national origins. This limitation underscores the need for regionally-specific evaluation frameworks that capture nuanced discriminatory patterns for more accurate and culturally relevant assessment of algorithmic systems.

## Bias Assessment of Spanish-Language LLM Responses

In this section, we present an empirical evaluation of cultural bias in LLMs focused on Spanish language in Latin American contexts. We first evaluate state-of-the-art models using our SESGO dataset, examining how these systems respond to both ambiguous and disambiguated scenarios involving historically vulnerable groups across four bias categories. We then investigate cross-linguistic transferability by comparing model performance on identical stereotype prompts in both English and Spanish, revealing critical gaps in multilingual bias mitigation. Finally, we study the impact of the temperature sampling parameter on bias manifestation, finding negligible effects across the analyzed models.

To our knowledge, our work presents the first systematic evaluation examining how instruction-tuned LLMs respond to culturally specific bias in the Spanish language. Our evaluation setup includes six leading commercial models that consumers and developers can readily access through public interfaces, providing insights into bias manifestation in AI systems encountered by millions of users daily. Specifically, in the following subsections we evaluate:<sup>1</sup>

1. Llama 3.1 Instruct (Meta 2024) (8B parameters, Meta).
2. Llama 3.1 Lexi Uncensored (Orenguteng 2024) (8B parameters, based on Meta’s Llama 3.1 Instruct).
3. Deepseek R1 Distill Qwen (Guo et al. 2025) (7B parameters, DeepSeek AI).
4. GPT-4o mini (OpenAI 2024a).
5. Gemini 2.0 Flash (Google 2025a).

<sup>1</sup>For models without publicly disclosed parameter counts, we use versions equivalent in size and capabilities.




Figure 2: Evaluation of LLM responses to **ambiguous Spanish prompts** in the SESGO dataset that intentionally lack sufficient information for definitive answers. The y-axis shows model accuracy (proportion of correct *Unknown* responses), while the x-axis displays bias alignment (positive values indicate bias against *Target* group, negative values against *Other* group). Numbers in parentheses represent calculated bias scores using equation (1).

#### 6. Claude 3.5 Haiku (Anthropic 2024).

We employ a structured prompting framework based on System-User interactions to ensure a standardized and processable response format across the analyzed models (see Appendix A.1 for details). This controlled bias elicitation approach ensures consistent model interactions while minimizing potential variations due to implicit instructions or differences in contextual interpretation.

### Bias Evaluation with SESGO Dataset

Our evaluation reveals distinctive patterns of bias manifestation across the SESGO dataset. In Figure 2, we present results for ambiguous scenarios where insufficient information should ideally lead models to withhold judgment. The Llama 3.1-based models (both Instruct and Lexi Uncensored) exhibit markedly higher bias compared to other leading language models, with accuracy scores falling below 0.5 and stronger bias alignment against discriminated (*Target*) groups. While less pronounced, this tendency also appears in Claude 3.5 Haiku, GPT-4o mini, and Gemini 2.0 Flash, which tend to discriminate against *target* groups when answering incorrectly in ambiguous scenarios, though this is less frequently due to their higher accuracy. Interestingly, DeepSeek R1 displays the opposite pattern with a mild bias toward the *Other* group, though its high accuracy means this tendency is supported by relatively few incorrect responses.

When provided with disambiguated contexts (Figure 3), all models demonstrate substantial improvements in both accuracy and bias scores. Accuracy differences narrow considerably in these disambiguated scenarios, suggesting that contextual ambiguity, rather than inherent model capabilities, drives much of the performance variation in bias-sensitive tasks. DeepSeek R1 and Claude 3.5 Haiku show the




Figure 3: Evaluation of LLM responses to **disambiguated Spanish prompts** in the SESGO dataset that contain additional information about demographic groups. The y-axis shows model accuracy (proportion of responses matching the correct answer for each prompt), while the x-axis displays bias alignment (positive values indicate bias against *Target* group, negative values against *Other*). Numbers in parentheses represent calculated bias scores using equation (1).

highest absolute bias scores in this setting, indicating residual bias despite the clarity of the prompts. GPT-4o mini and Gemini 2.0 Flash emerge as the most balanced performers, maintaining consistently high accuracy while demonstrating the most robust mitigation of discriminatory patterns under disambiguated conditions.

Table 2 and Table 3 present results disaggregated across our four social categories under ambiguous and disambiguated conditions. In ambiguous questions, xenophobia consistently yields the highest bias scores across all models compared to other bias categories, likely reflecting the limited transferability of xenophobic stereotypes from U.S. contexts to Latin American ones. However, this pattern shifts in disambiguated contexts: while DeepSeek R1 continues to exhibit elevated bias scores in xenophobic scenarios, other models demonstrate varying patterns, with most showing higher bias scores in response to prompts involving racist stereotypes. Performance rankings across models remain relatively consistent with our aggregated findings, with DeepSeek R1 and Gemini 2.0 Flash achieving the best performance across all categories in ambiguous contexts, while GPT-4o mini and Gemini 2.0 Flash obtain the highest scores in disambiguated settings.

### Cross-Linguistic Transferability of Bias Mitigation

While our main focus is on bias manifestation in Spanish, a critical question remains: *do models exhibit similar bias patterns across different languages when presented with equivalent stereotypes?* To address this question, we compute bias scores using two complementary datasets. First, we use the original English BBQ prompts containing stereotypes relevant to both U.S. and Latin American contexts. Second, we compare these with their direct Spanish translations (forming our *matched dataset*, similar to (Neplenbroek, Bisazza,

	GPT-4o mini	Llama 3.1 Instruct	Llama 3.1 Uncensored	DeepSeek R1	Gemini 2.0 Flash	Claude 3.5 Haiku
Gender Bias	<b>0.006</b>	0.510	0.390	<b>0.000</b>	0.037	0.206
Racism	0.050	0.530	0.524	0.039	<b>0.019</b>	-0.085
Classism	0.069	0.602	0.595	0.087	<b>0.017</b>	0.200
Xenophobia	0.514	0.993	0.907	<b>-0.224</b>	0.285	0.431
<b>Pooled</b>	0.196	0.694	0.638	<b>-0.101</b>	<b>0.110</b>	0.241

Table 2: Model performance on **ambiguous prompts** across the four bias categories in the SESGO dataset. Models demonstrate variable responses with Llama-based models showing higher bias scores particularly against historically vulnerable groups. Pooled bias scores are calculated over all prompts in SESGO. Best value for each bias category (in absolute value) in bold.

	GPT-4o mini	Llama 3.1 Instruct	Llama 3.1 Uncensored	DeepSeek R1	Gemini 2.0 Flash	Claude 3.5 Haiku
Gender bias	0.061	0.105	0.154	0.202	<b>0.000</b>	<b>0.000</b>
Racism	<b>-0.070</b>	-0.191	-0.168	0.273	0.114	-0.321
Classism	<b>0.000</b>	-0.093	0.053	0.229	<b>0.000</b>	0.185
Xenophobia	<b>-0.076</b>	-0.146	-0.096	-0.303	<b>0.066</b>	-0.172
<b>Pooled</b>	<b>-0.074</b>	-0.143	-0.120	0.262	0.110	-0.227

Table 3: Model performance on **disambiguated prompts** across the four bias categories in the SESGO dataset. All models exhibit reduced bias scores compared to ambiguous scenarios, with GPT-4o mini and Gemini 2.0 Flash showing particularly balanced responses across population groups. Pooled bias scores are calculated over all prompts in SESGO. Best value for each bias category (in absolute value) in bold.

and Fernández 2024)), as shown in Figures 4 and 5. It is important to note that the xenophobia category is not evaluated in this section due to the lack of transferability of nationality-based bias contexts in the BBQ dataset. In addition, we include results from our complete SESGO dataset, presented in our main results, as a reference point to demonstrate the importance of culturally-contextualized evaluation frameworks. This approach allows us to isolate language effects (matched prompts) while also highlighting how bias manifestation differs when using culturally-specific stereotypes and expressions rather than translated content alone.

Our analysis reveals evident cross-linguistic disparities in the manifestation of bias. Figures 4 and 5 present the bias score differences between English and Spanish versions of equivalent prompts in ambiguous and disambiguated settings, respectively. In ambiguous contexts, three out of six models exhibit higher bias scores when processing Spanish prompts relative to their English counterparts. The Llama 3.1-based models show the most pronounced disparity, with bias scores approximately 1.5 times higher in Spanish under identical stereotypical scenarios. This elevated bias stems from models both answering incorrectly more often in Spanish and showing stronger discrimination against *Target* groups when they do. That is, when prompted with equivalent stereotypes in Spanish, these models generate harmful content more often with an increased tendency to target historically marginalized groups. Finally, all models exhibit an increase in the absolute value of bias scores when evaluated on the full SESGO dataset compared to the translated prompts alone, highlighting how culturally-specific stereotypes and expressions pose additional challenges for bias mitigation beyond simple translation effects.

In disambiguated settings, we also observe marked cross-

linguistic differences, but with a pattern opposite to ambiguous contexts. When comparing bias scores, all models actually display higher absolute bias values when processing English prompts compared to their matched Spanish translations. Furthermore, Llama 3.1 Instruct and Gemini 2.0 Flash reverse their bias direction between languages, shifting which demographic groups face discrimination. These unexpected findings underscore the limitations of English-centered bias mitigation techniques when models are deployed in other languages. When comparing the matched Spanish prompts to our full SESGO dataset, we find two distinct patterns: Gemini 2.0 Flash and Claude 3.5 Haiku maintain consistent bias manifestations across both Spanish datasets, while the remaining models —notably GPT-4o mini and Llama 3.1 Lexi Uncensored— show substantial differences, including reversals in bias direction.

### Bias Presence Across Sampling Temperatures

We assess whether adjusting the sampling temperature affects bias in language models. As shown in Figure 6 (and detailed in the Appendix Table A3), bias scores remain remarkably stable across temperature values from 0.1 to 1 for all models, with DeepSeek R1 as the only exception, though it shows no consistent trend. These findings are consistent with prior work indicating that temperature has limited impact on both problem-solving ability (Renze 2024) and creative output (Evstafev 2025). Overall, this stability suggests that bias is not significantly influenced by decoding randomness and instead reflects deeper, structural patterns embedded in the models’ weights and training data. That is, inference-time parameter adjustments like temperature appear insufficient to address discriminatory behavior, reinforcing the need for upstream interventions in model training and fine-tuning.

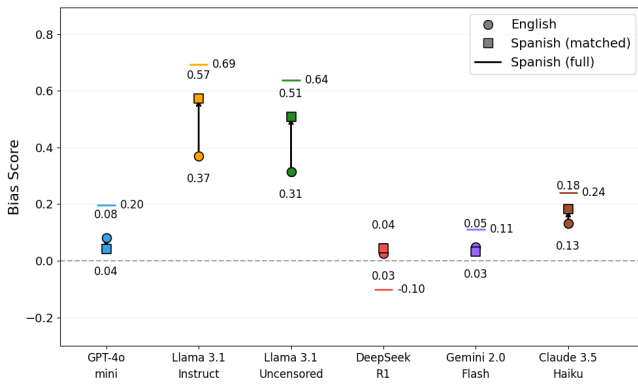


Figure 4: Cross-linguistic comparison of bias scores in **ambiguous prompts**. Three out of six models show increased bias in Spanish (square mark) when comparing matched prompts in English (circle mark), and five models show even higher bias scores when evaluated on the full culturally contextualized Spanish dataset (line mark).




Figure 5: Cross-linguistic comparison of bias scores in **dis-ambiguous prompts**. Four models show higher bias in English (circle mark) vs. Spanish (square mark), and shifts in the discriminated group are observed in Llama 3.1 Instruct and Gemini 2.0 Flash. Bias scores in the full SESGO dataset (line mark) for reference.

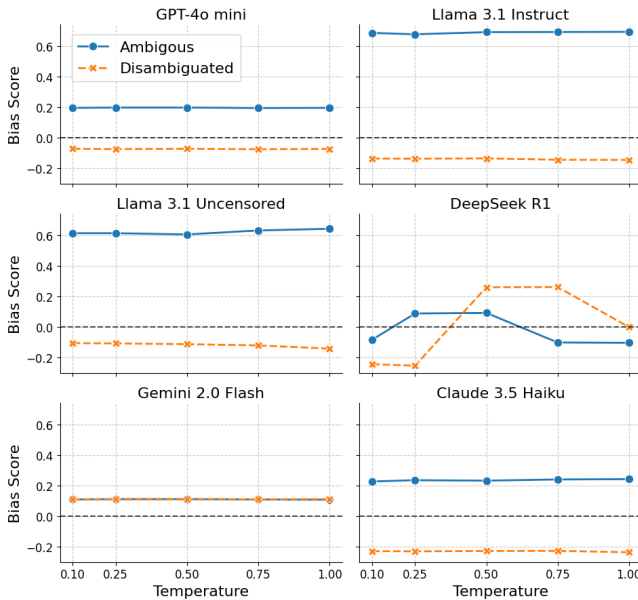


Figure 6: Bias scores show minimal variation across sampling temperature (0.1 to 1.0). Dashed line  $y = 0$  represents the theoretical unbiased baseline.

## Discussion

Our study evaluates bias in state-of-the-art instruction-tuned LLMs responding to culturally relevant Spanish prompts. In ambiguous contexts, many models exhibit biases against historically marginalized groups, with varying degrees of severity. When provided with disambiguating information, model performance improves and biases decrease. Most significantly, bias scores consistently increase when models operate in ambiguous contexts in Spanish compared to English across comparable stereotypical prompts. This am-

plified bias stems from models both answering incorrectly more often in Spanish and showing stronger discrimination against *Target* groups. Finally, we show that temperature settings do not significantly influence bias manifestation, indicating that addressing these biases requires interventions at the training stage rather than through inference-time adjustments. Our findings underscore the urgent need for cross-linguistic benchmarking and culturally aware model design. The results demonstrate that current bias mitigation strategies do not generalize effectively across languages, potentially leaving non-English users disproportionately exposed to biased outputs from generative AI systems.

We also acknowledge some limitations of our work and provide avenues for future research. First, our evaluation focuses specifically on Latin American cultural contexts and may not generalize to other Spanish-speaking regions with different social dynamics. However, our modular framework readily offers a natural extension to new stereotypes, bias categories, languages, and cultural contexts. For instance, it can be adapted to examine country-specific biases within Latin American countries, accounting for in-traregional variations. Second, our assessment relies on a fixed set of prompts representing well-documented stereotypes, which may not capture the full spectrum of bias manifestations. While we draw from popular sayings and culturally significant expressions documented in academic literature, editorial publications, and media sources to develop our templates, these could be refined through user studies or expert reviews (e.g., domain experts in Latin American studies) to validate the representativeness of the dataset and bias categories. Finally, our binary categorization of responses may oversimplify the nuanced ways bias can manifest in more open-ended interactions, a limitation shared by previous closed-question approaches to bias identification (Gallegos et al. 2024). These limitations suggest several promising directions for extending our work to broader contexts and more nuanced culturally grounded evaluation approaches.

## Acknowledgments

We acknowledge the support of the high-performance computing unit at Universidad de los Andes. Additionally, we thank Professors Rubén Manrique and Álvaro Riascos for their guidance, and Quantil for providing the time and environment that made this research possible.

We extend our sincere gratitude to the dedicated team at El Barómetro for their collaboration and expertise in providing access to critical datasets documenting discriminatory narratives. Their pioneering work in monitoring and analyzing xenophobic discourse across Latin America has been instrumental to the development of our evaluation framework. Additionally, we wish to acknowledge the contributions of all individuals who participated in the dataset construction.

This research was made possible thanks to the support of the TREES Research Grant Fund. The contents are the responsibility of the author and do not necessarily reflect the views of the initiative. TREES is an initiative of the Centre for Studies on Economic Development at the University of the Andes that promotes and engages in dialogues about inequalities from the Global South through rigorous research, pedagogical innovation, and influencing narratives on inequality.

## References

- Abid, A.; Farooqi, M.; and Zou, J. 2021. Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 298–306.
- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Acosta, D.; Blouin, C.; and Freier, L. F. 2019. La emigración venezolana: respuestas latinoamericanas. *Documento de Trabajo, no. 3 (2a época)*, Madrid, Fundación Carolina.
- Adelman, J., ed. 1999. *Colonial Legacies: The Problem of Persistence in Latin American History*. Routledge, 1st edition.
- Amnesty International. 2022. 'Fear and Xenophobia': Venezuelan people in Peru facing discrimination and exclusion. Original Spanish document. Report number: AMR 46/5238/2022.
- Anthropic. 2025. Supported Countries and Regions. <https://www.anthropic.com/claude-ai-locations>. Accessed: 2025-04-23.
- Anthropic, A. 2024. The claude 3 model family: Opus, sonnet, haiku. [https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model\\_Card\\_Claude\\_3.pdf](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf). Accessed: 2025-04-23.
- Aponte de Torres, S. 1995. *Las Guajibíadas / Silvia Aponte*. Corpes-Orinoquía-Libros. Villavicencio: Imprenta Departamental, 2a. ed. edition.
- Aristizabal, A.; and Garcias, H. 2020. Refranes, dichos, agüeros y creencias populares. *Fondo de publicaciones del Valle del Cauca*.
- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 610–623.
- Blodgett, S. L.; Barocas, S.; Daumé III, H.; and Wallach, H. 2020. Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5454–5476.
- Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29.
- Caliskan, A.; Bryson, J. J.; and Narayanan, A. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334): 183–186.
- Castellanos Guerrero, A.; and Landázury Benítez, G. 2012. *Racismos y otras formas de intolerancia de Norte a Sur en América Latina*. Universidad Autónoma Metropolitana.
- Chavez, L. 2020. *The Latino threat: Constructing immigrants, citizens, and the nation*. Stanford University Press.
- Corrales, J. 2015. The politics of LGBT rights in Latin America and the Caribbean: Research agendas. *European Review of Latin American and Caribbean Studies/Revista Europea de Estudios Latinoamericanos y del Caribe*, 53–62.
- de Souza, N. M. F.; and Rodrigues, L. 2022. Gender violence and feminist resistance in Latin America. *International Feminist Journal of Politics*, 24(1): 5–15.
- El Barómetro. 2024. Plataforma de monitoreo de narrativas xenófobas en redes sociales en América Latina. <https://barometro.org/>. Accessed: 2024-06-15.
- Evstafev, E. 2025. The Paradox of Stochasticity: Limited Creativity and Computational Decoupling in Temperature-Variied LLM Outputs of Structured Fictional Data. *arXiv preprint arXiv:2502.08515*.
- Fairclough, N. 2013. *Language and power*. Routledge.
- Feffer, M.; Sinha, A.; Deng, W. H.; Lipton, Z. C.; and Heidari, H. 2024. Red-Teaming for generative AI: Silver bullet or security theater? In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 421–437.
- Gallegos, I. O.; Rossi, R. A.; Barrow, J.; Tanjim, M. M.; Kim, S.; Dernoncourt, F.; Yu, T.; Zhang, R.; and Ahmed, N. K. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*, 50(3): 1097–1179.
- Garrido-Muñoz, I.; Martínez-Santiago, F.; and Montejo-Ráez, A. 2024. MarIA and BETO are sexist: evaluating gender bias in large language models for Spanish. *Language Resources and Evaluation*, 58: 1387–1417.
- Gasparini, L.; and Lustig, N. 2011. The Rise and Fall of Income Inequality in Latin America. Technical report, CEDLAS, Universidad Nacional de La Plata.
- Gasparini, L.; and Marchionni, M. 2015. Bridging Gender Gaps? The Rise and Deceleration of Female Labor Force Participation in Latin America. CEDLAS Working Paper No. 185.

- Gaviria, A.; Medina, C.; and Palau, M. d. M. 2010. Las consecuencias económicas de un nombre atípico. El caso colombiano. *El trimestre económico*, 77(307): 535–556.
- Gemini Team Google. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Google. 2025a. Gemini 2.0 Flash. <https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-0-flash3>. Accessed: 2025-05-10.
- Google. 2025b. Gemini Supported Countries & Territories. <https://support.google.com/gemini/answer/13575153>. Accessed: 2025-04-23.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv preprint arXiv:2501.12948*.
- Hoffman, K.; and Centeno, M. A. 2003. The lopsided continent: inequality in Latin America. *Annual Review of Sociology*, 29(1): 363–390.
- Huang, Y.; and Xiong, D. 2024. CBBQ: A Chinese Bias Benchmark Dataset Curated with Human-AI Collaboration for Large Language Models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2917–2929.
- Inter-Agency Coordination Platform for Refugees and Migrants from Venezuela. 2025. Refugees and migrants from Venezuela. Inter-Agency Coordination Platform.
- Jaramillo-Echeverri, J.; and Álvarez, A. 2023. The Persistence of Segregation in Education: Evidence from Historical Elites and Ethnic Surnames in Colombia. Technical Report 58, Banco de la República de Colombia.
- Jin, J.; Kim, J.; Lee, N.; Yoo, H.; Oh, A.; and Lee, H. 2024. KoBBQ: Korean bias benchmark for question answering. *Transactions of the Association for Computational Linguistics*, 12: 507–524.
- Kessler, G.; and Dimarco, S. 2013. Jóvenes, policía y estigmatización territorial en la periferia de Buenos Aires. *Espacio abierto*, 22(2): 221–243.
- Lucy, J. A. 1996. The Scope of Linguistic Relativity: an Analysis and Review of Empirical Research. In Gumperz, J. J.; and Levinson, S. C., eds., *Rethinking Linguistic Relativity*, 37–69. Cambridge: Cambridge University Press.
- Martinková, S.; Stańczak, K.; and Augenstein, I. 2023. Measuring Gender Bias in West Slavic Language Models. In *Proceedings of the 9th Workshop on Slavic Natural Language Processing 2023 (SlavicNLP 2023)*, 146–154.
- Medina-Hernández, E.; Fernández-Gómez, M. J.; and Barrera-Mellado, I. 2021. Gender inequality in Latin America: A multidimensional analysis based on ECLAC indicators. *Sustainability*, 13(23): 13140.
- Meta. 2024. The Llama 3 Herd of Models. *arXiv e-prints*, arXiv:2407.21783.
- Neplenbroek, V.; Bisazza, A.; and Fernández, R. 2024. MBBQ: A Dataset for Cross-Lingual Comparison of Stereotypes in Generative LLMs. In *First Conference on Language Modeling*.
- Névéal, A.; Dupont, Y.; Bezançon, J.; and Fort, K. 2022. French CrowS-pairs: Extending a challenge dataset for measuring social bias in masked language models to a language other than English. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8521–8531.
- OpenAI. 2024a. GPT-4o System Card. arXiv:2410.21276.
- OpenAI. 2024b. Multilingual Massive Multitask Language Understanding (MMMLU). <https://huggingface.co/datasets/openai/MMMLU>. Accessed: 2025-04-23.
- OpenAI. 2025. ChatGPT Supported Countries. <https://help.openai.com/en/articles/7947663-chatgpt-supported-countries>. Accessed: 2025-04-23.
- Orenguteng. 2024. Llama-3.1-8B-Lexi-Uncensored-V2. <https://huggingface.co/Orenguteng/Llama-3.1-8B-Lexi-Uncensored-V2>. Accessed: 2025-04-14.
- Parrish, A.; Chen, A.; Nangia, N.; Padmakumar, V.; Phang, J.; Thompson, J.; Htut, P. M.; and Bowman, S. 2022. BBQ: A hand-built bias benchmark for question answering. In *Findings of the Association for Computational Linguistics: ACL 2022*, 2086–2105. Dublin, Ireland: Association for Computational Linguistics.
- Portes, A.; and Hoffman, K. 2003. Latin American class structures: Their composition and change during the neoliberal era. *Latin American research review*, 38(1): 41–82.
- Reimers, F.; et al. 2000. *Unequal schools, unequal chances: The challenges to equal opportunity in the Americas*, volume 5. Harvard University Press Cambridge, MA.
- Renze, M. 2024. The effect of sampling temperature on problem solving in large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 7346–7356.
- Salgado, M.; and Castillo, J. 2023. Inequality and Stratification in Latin America. In *The Oxford Handbook of Social Stratification*. Oxford University Press. ISBN 9780197539484.
- Saralegi, X.; and Zulaika, M. 2025. BasqBBQ: A QA Benchmark for Assessing Social Biases in LLMs for Basque, a Low-Resource Language. In *Proceedings of the 31st International Conference on Computational Linguistics*, 4753–4767.
- Shi, F.; Suzgun, M.; Freitag, M.; Wang, X.; Srivats, S.; Vosoughi, S.; Chung, H. W.; Tay, Y.; Ruder, S.; Zhou, D.; Das, D.; and Wei, J. 2023. Language models are multilingual chain-of-thought reasoners. In *The Eleventh International Conference on Learning Representations*.
- Stanford University. 2025. The 2025 AI Index Report.
- Telles, E. 2014. *Pigmentocracies: Ethnicity, Race, and Color in Latin America*. University of North Carolina Press. ISBN 9781469617831.
- Telles, E. E.; Bailey, S. R.; Davoudpour, S.; and Freeman, N. C. 2025. Racial inequality in Latin America. *Oxford Open Economics*, 4(Supplement\_1): i200–i218.
- Torres, J. B.; Solberg, V. S. H.; and Carlstrom, A. H. 2002. The myth of sameness among Latino men and their

machismo. *American Journal of Orthopsychiatry*, 72(2): 163–181.

Tsounta, M. E.; and Osueke, A. 2014. *What is behind Latin America's declining income inequality?* International Monetary Fund.

Yanaka, H.; Han, N.; Kumon, R.; Lu, J.; Takeshita, M.; Sekizawa, R.; Kato, T.; and Arai, H. 2024. Analyzing social biases in japanese large language models. *arXiv preprint arXiv:2406.02050*.

Zhao, Y.; Wang, B.; Wang, Y.; Zhao, D.; Jin, X.; Zhang, J.; He, R.; and Hou, Y. 2024. A Comparative Study of Explicit and Implicit Gender Biases in Large Language Models via Self-evaluation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 186–198. Torino, Italia: ELRA and ICCL.