

MEASURING OPINIONATED COMMUNICATION

SIMULATING CITIZENS, AMPLIFYING STEREOTYPES?

Can AI really act like a typical citizen?

Ana Jovančević, Ljubiša Bojić, Olga Zagovora, & Velibor Ilić

INTRODUCTION

Large Language Models (LLMs) are increasingly used to simulate citizens in research and policy. We examined whether AI models reproduce human attitudes on gender equality and immigration across Ireland, the UK, Germany and Serbia.

METHODS

6 LLMs: GPT-5, Grok, Gemini, Gemma, Mistral and DeepSeek;
4 countries (Ireland, UK, Germany, and Serbia);
2 prompt types (Neutral vs Persona);
Topics: Gender Equality & Immigration
Compared with European Social Survey data (Wave 11).

STUDY DESIGN

Neutral Prompt → LLM Response

Persona Prompt (Typical Citizen) → LLM Response

Comparison with real citizens.

MAIN FINDINGS

1. PERSONA EFFECTS

LLMs became more conservative when answering as a typical citizen.


2. COUNTRY EFFECTS

Results varied by language and country, with stronger shifts in Serbian prompts.


3. HUMAN COMPARISON

LLMs moved toward human attitudes but did not fully match real citizens.

How similar are LLM responses on Immigration attitudes with human responses?



How similar are LLM responses on Gender Equality attitudes with human responses?



WHY IT MATTERS?

AI-generated citizens are increasingly used for social simulations. Persona prompting can amplify stereotypes and may distort conclusions if treated as a substitute for real public opinion.

TAKE-HOME MESSAGE: AI can imitate societies, but it can also distort them. Language, persona and model choice significantly affect the opinions AI expresses.