Media and Communication
2025 • Volume 13 • Article 9677
https://doi.org/10.17645/mac.9677

ARTICLE

Open Access Journal

Ideology and Policy Preferences in Synthetic Data: The Potential of LLMs for Public Opinion Analysis

Keyeun Lee [1], Jaehyuk Park [2], Suh-hee Choi [3], and Changkeun Lee [2]

[1] Department of Communication, Seoul National University, South Korea
[2] KDI School of Public Policy and Management, South Korea
[3] Department of Geography, Kyung Hee University, South Korea

Correspondence: Changkeun Lee (cklee@kdischool.ac.kr)

Submitted: 21 November 2024 **Accepted:** 5 March 2025 **Published:** 22 May 2025

Issue: This article is part of the issue "AI, Media, and People: The Changing Landscape of User Experiences and Behaviors" edited by Jeong-Nam Kim (University of Oklahoma) and Jaemin Jung (Korea Advanced Institute of Science and Technology), fully open access at https://doi.org/10.17645/mac.i475

Abstract

This study investigates whether large language models (LLMs) can meaningfully extend or generate synthetic public opinion survey data on labor policy issues in South Korea. Unlike prior work conducted on people's general sociocultural values or specific political topics such as voting intentions, our research examines policy preferences on tangible social and economic topics, offering deeper insights for news media and data analysts. In two key applications, we first explore whether LLMs can predict public sentiment on emerging or rapidly evolving issues using existing survey data. We then assess how LLMs generate synthetic datasets resembling real-world survey distributions. Our findings reveal that while LLMs capture demographic and ideological traits with reasonable accuracy, they tend to overemphasize ideological orientation for politically charged topics—a bias that is more pronounced in fully synthetic data, raising concerns about perpetuating societal stereotypes. Despite these challenges, LLMs hold promise for enhancing data-driven journalism and policy research, particularly in polarized societies. We call for further study into how LLM-based predictions align with human responses in diverse sociopolitical settings, alongside improved tools and guidelines to mitigate embedded biases.

Keywords

AI-generated text; ChatGPT; large language models; news media; policy preferences; public opinions

1. Introduction

With the advancement of large language models (LLMs), there has been growing interest in how well AI-generated factual and opinionated text can resemble human output among academicians and journalists.

AI-generated factual text has begun affecting routine human tasks such as combining and summarizing information (Glickman & Zhang, 2024). A recent focus of academic studies also involves how LLMs can generate text that mirrors the content, style, tone, and grammatical traits of specific demographic groups (Argyle et al., 2023; Gerosa et al., 2024; Harding et al., 2024; Sun et al., 2024). Beyond academic circles, acknowledging the benefits in terms of efficiency and timeliness, journalists are also testing LLMs to gauge public opinion on policy issues. Traditionally, news media have relied on opinion surveys to track public sentiment and shape discourse by presenting "snapshots" of crucial social, political, and economic issues. However, such conventional methods face challenges—they are costly, time-consuming, and often lag behind rapidly changing events. Fast, cost-efficient AI-based approaches are growing in appeal, especially as media outlets strive to capture evolving public attitudes more promptly. For example, *The Atlantic*'s tech reporter experimented by prompting ChatGPT to act as different archetypal voters ("MAGA zealots," "suburban moms," etc.), discovering that a "40-year-old conservative man from rural Ohio" persona produced "vividly partisan" rhetoric (Desai, 2023). This demonstrates how easily the AI can mirror extreme opinions when instructed to adopt a specific viewpoint.

Meanwhile, some news media are considering AI's capacity to generate focus-group-style responses to policy proposals. So far, most experiments remain in a testing phase—not yet having formal "AI-generated polls" in mainstream outlets—and reputable publications like *The New York Times* and BBC have announced cautionary guidelines. The BBC, for instance, explicitly forbids using LLMs (BBC, 2025). Still, the intrigue surrounding AI-based public opinion simulations persists, underscoring the tension between editorial caution and innovative ambition.

Despite extensive scholarly warnings and recent journalism guidelines urging a careful approach, global news outlets increasingly rely on LLMs to produce political comparisons and predictive analyses. For instance, in South Korea, the conservative newspaper *Chosun Ilbo* published an article asking ChatGPT about North Korean leader Kim Jong Un (G. Y. Kim, 2023), while the progressive newspaper *Hankyoreh* published a column inquiring about potential job destruction caused by LLMs (S. Lee, 2023). Such rapidly produced, AI-crafted stories appealed to readers with strong political views—whether supportive or critical of particular leadership styles—yet often lack verifiable data or current context. Scholars emphasize that LLMs cannot replace expert analysis or rigorous reporting; nonetheless, the allure of fast, eye-catching AI-driven content widens the gap between recommended guidelines and actual newsroom practice.

This article examines how LLMs can complement traditional opinion surveys in South Korea, where LLMs have already gained wide traction, particularly on polarized social and political issues. We begin by considering a scenario in which a new topic—or a fresh angle on an existing one—demands timely insights. In this setting, we test whether LLMs, provided individual-level survey data, can effectively predict public sentiment toward that emerging issue. We then turn to whether LLMs can generate synthetic data that closely reflect real-world survey distributions, thus potentially aiding researchers and journalists in predicting trends and outcomes with greater detail and at a lower cost.

In doing so, we pay particular attention to labor policy, where South Korea's left- and right-leaning groups hold notably divergent views. While minimum wage controversies are highly politicized, extending the retirement age typically engenders far less ideological friction. Recent studies such as Rozado (2024) caution that LLMs may carry political biases, but few have examined how these biases vary between politicized and non-politicized questions. Building on the literature on ideological polarization (Caughey et al., 2019; McCarty et al., 2016), we investigate whether LLM-driven responses amplify partisan rifts on hot-button labor issues like the minimum wage compared to relatively neutral topics such as retirement age. Our preliminary findings reveal that AI-driven outputs can overestimate the effect of users' ideological orientation, underscoring the need for vigilance when employing LLMs to extend or create new data. By highlighting these potential pitfalls, we aim to offer guidance for researchers and media outlets considering AI-based methods to gauge public sentiment—particularly in highly polarized environments.

This article contributes to the literature in several ways. First, unlike many studies that center on general sociocultural topics, including people's values, or specific political perceptual or behavioral topics, such as voting intentions, we focus on specific policy issues. By examining how LLMs capture—or misrepresent—public preferences on concrete policies, our research supports a shift toward policy-oriented journalism rather than surface-level partisan divides.

Second, we situate our analysis within the South Korean context, where political polarization has intensified even though actual differences in public policy preferences often remain modest (Cheong & Haggard, 2023). While political orientation is generally presumed to be a strong predictor of policy preferences, this assumption warrants re-examination in South Korea, where ideological identities dominate public discourse but may not map neatly onto specific policy positions. This fills a critical gap in the predominantly North American- and European-focused literature, providing a fresh perspective on the complex relationship between political identity and policy choices.

Finally, the rising prevalence of AI- and human-generated text in the media has already altered how people consume information (Yang & Menczer, 2024). Some AI-generated content has been flagged as malicious, raising concerns about its distorting effects on the data ecosystem. Blending opinionated comments and news text can confuse audiences—particularly older generations who may be less familiar with AI technologies (Moravec et al., 2024). Furthermore, the problem of AI "hallucination," where models produce inaccurate or misleading content, underscores the need for better training data, greater transparency, and continued ethical oversight (Gerosa et al., 2024; Patel, 2024). We contribute to the extant body of knowledge by showing how these challenges also apply to public opinion data, reinforcing the importance of careful validation and critical scrutiny when using LLM-generated data in this field.

2. Literature Review

2.1. AI in Journalism: Opportunities and Risks

Recent scholarship underscores both the significant promise and potential pitfalls of integrating advanced language models into newsrooms. Caswell (2024) shows how media organizations have moved from simpler automated outputs, such as basic data-driven stories, to generative AI, which has resulted in marked gains in productivity and the possibility of targeted, audience-specific content. However, the author also urges the creation of dedicated "AI editors" who can identify bias and maintain rigorous journalistic standards, pointing out that traditional gatekeeping roles might weaken if automation is not guided by specialized oversight. Pan et al. (2023) highlight the risk of widespread "misinformation pollution," as these models can generate convincing but false or biased statements in large volumes, often surpassing both human and automated fact-checkers. Fletcher and Nielsen (2024) report public skepticism toward AI-driven stories, with fewer than one in ten people willing to pay extra for content written by machines. Other researchers also doubt if AI systems can consistently meet professional editorial benchmarks, expressing worries about accuracy and transparency. (Binz & Schulz, 2023; Brigham et al., 2024). Brigham et al. (2024) reveal ethical concerns, such as journalists inadvertently exposing confidential or copyrighted materials to LLMs, and describe minimal human revision of AI drafts, which can let factual errors slip through.

Collectively, these findings emphasize that LLMs can streamline production and even personalize news for different readers, but they also risk bias, misinformation, and a gradual decline in editorial integrity when used uncritically. For our study, these dynamics underscore the importance of robust oversight, ethical safeguards, and editorial clarity regarding AI-generated content—principles that likewise apply to the use of LLMs in public opinion research.

2.2. Using AI to Extend Public Opinion Data

Advances in LLMs have stimulated a surge of research on whether AI-generated text can effectively replicate or extend public opinion data. Some applications focus on AI-driven polling, where synthetic survey responses can reduce costs and time relative to conventional polls (Berger et al., 2024). While these pilot initiatives show promise in forecasting policy preferences and voter behavior, they also amplify longstanding concerns about algorithmic bias and the reliability of training data (Berger et al., 2024; Kennedy et al., 2022). Other studies build on this foundation by examining the capacity of LLMs to not only generate textual responses but also to capture more nuanced socioeconomic and political expressions (Argyle et al., 2023; Feng et al., 2023).

A key focus of recent work is conditioning the language model with demographic or ideological cues. Researchers have shown that ChatGPT-based systems can yield textual outputs resembling specific subgroups' ways of expressing political or socioeconomic attitudes (Amirova et al., 2024; Argyle et al., 2023). This approach purports to capture the probabilistic distributions of real human responses under given demographic or ideological constraints. Proponents argue that such "silicon subjects" (Argyle et al., 2023; Sun et al., 2024) offer an avenue to explore large-scale, fine-grained variations in opinion without incurring the high costs and potential sampling biases of traditional surveys (Gerosa et al., 2024; Hutson, 2023; Pachot & Petit, 2024).

Despite these potential benefits, several researchers emphasize the concept of algorithmic fidelity, the extent to which an LLM accurately reflects the attitudes and opinions of a targeted human subgroup (Amirova et al., 2024; Argyle et al., 2023). The premise is that if sufficiently detailed conditioning prompts are provided (covering age, gender, ideology, socioeconomic status, and so forth) the generated text will probabilistically mirror human responses in real-world settings. However, as Amirova et al. (2024) note, there can be inconsistencies or discrepancies between what the model produces and actual subgroup attitudes, especially if the model's training data lack diversity or reflect outdated cultural contexts. A related challenge is persona simulation, where researchers prompt the model to "act as if" it possesses certain cognitive limitations or ideological commitments (Aher et al., 2023; Gerosa et al., 2024; Kotek et al., 2023; Milička et al., 2024). While these techniques may yield lifelike responses, they also risk producing hyper-accuracy distortion (Amirova et al., 2024) or "correct answer" effects (Park et al., 2024). In some cases, the model's factual knowledge overrides attempts to simulate uncertainty or misinformation, resulting in text that is too well-informed and thus unrepresentative of real human respondents.

A crucial limitation arises from sociopolitical biases embedded in LLMs (Aher et al., 2023; Feng et al., 2023; Kotek et al., 2023; Park et al., 2024). Because models train on data that may overrepresent certain demographic groups or partisan content, they often exhibit skewed outcomes when asked to mimic diverse populations (Aher et al., 2023). Researchers highlight issues such as gender bias (Kotek et al., 2023) and the tendency to produce uniformly "safe" responses that do not capture the full range of ideologically extreme viewpoints (Park et al., 2024). Dillion et al. (2023) and Hutson (2023) note that LLMs cannot replicate the intricate psychological or social processes underlying real human behavior, such as lying, changing opinions over time, or experiencing moral dilemmas. Some studies also observe distortions in the model's adherence to specified personas or instructions. Milička et al. (2024) document instances where advanced ChatGPT-based systems fail to limit or "downplay" their own cognitive abilities when prompted with less-informed personas, thus providing responses that are more coherent and factual than a human subject with similar constraints might. In other words, even a well-tuned model may inadvertently slip into more knowledgeable or accurate modes of response, invalidating efforts to simulate ignorance, confusion, or bias.

Recent literature thus presents a mixed picture: On the one hand, LLMs open up avenues for low-cost, high-volume simulations of public opinion (Argyle et al., 2023; Gerosa et al., 2024); on the other hand, they introduce new methodological and ethical complications. Issues like overly uniform outputs, underestimation of extreme viewpoints, and reliance on training data that reflect outdated or biased cultural contexts limit the reliability of LLM-generated "silicon subjects" (Harding et al., 2024; Park et al., 2024). Harding et al. (2024) argue that language models cannot easily adapt to the rapid cultural or social shifts occurring in real human populations—especially as new moral standards and political events reshape beliefs in ways that may not be reflected in static training corpora. From the perspective of policy analysts, political scientists, and media organizations, AI-based polling and ChatGPT simulations remain attractive for exploring emergent trends or investigating hypothetical scenarios (Sun et al., 2024). However, to fully harness these tools, researchers must actively mitigate biases, continually validate model outputs against real-world data, and disclose the AI's role in the generation process. In doing so, they may foster more nuanced and trustworthy insights while recognizing that LLMs cannot yet replace the complexity and variability inherent in genuine human attitudes (Dillion et al., 2023; Hutson, 2023).

2.3. The South Korean Policy and Media Context

An important underlying assumption in comparing the results from human responses with those from LLM-generated responses is that a refined, particular demographic group with a specific expression of political positioning presents relatively stable attitudes toward social issues such as population policies (Han & Ding, 2024). It is achieved by LLM's connection of the defined profile to the text generation coherent with the trained information (Gerosa et al., 2024).

As most previous studies have been conducted in English-speaking or Western societies, especially the US, where political ideological identities and those of political parties have been relatively stabilized, the input of variables associated with political parties, ideological inclinations, and sociodemographic characteristics can more readily predict individuals' political positions (Argyle et al., 2023). Despite some common notions of the spectrum of left and right, totalitarian–authoritarian, and libertarian, variation in the way such terms are understood and utilized may occur because the conditions that shape the spectrum consist of more than one criterion and because of region-specific contexts (Gindler, 2021). Similarly, research on US political orientations suggests relatively consistent liberal–conservative divides, reflecting distinct moral intuitions and responses to uncertainty (K. R. Kim & Kang, 2013).

In contrast, South Korea's party system is characterized by less clearly demarcated ideological boundaries, weakening the representational ties between specific social groups and political organizations (Cheong & Haggard, 2023; Cho et al., 2019). Parties often shift their policy positions or alliances and do not consistently anchor themselves to a stable ideological identity (Cheong & Haggard, 2023). Nevertheless, a trend of polarization has gained traction in recent years, with supporters of the left party embracing more liberal stances and right-leaning constituents identifying as more conservative (Cheong & Haggard, 2023). One illustrative example is the minimum wage debate, which has been highly politicized in South Korea. Although the left typically presents minimum wage hikes as a remedy for inequality and a means to protect vulnerable workers, the right emphasizes concerns about potential burdens on small and medium-sized enterprises, arguing for greater labor market flexibility instead. Such divergences demonstrate how certain issues—like the minimum wage—can become rallying points for entrenched beliefs that map onto left and right orientations, even in a context where overall party identities remain fluid.

Given this complex ideological landscape, assumptions about the stability and predictability of political attitudes in South Korean contexts may not hold to the same extent as in the US or other Western settings. For LLMs, generating coherent profiles based on South Korean demographic and ideological variables can, therefore, be more challenging, potentially yielding more variance in simulated opinions (Cheong & Haggard, 2023). This highlights why cross-national studies of LLM-generated responses need to consider regional political traditions and the strength (or weakness) of partisan identities when interpreting outcomes—particularly on politicized topics such as the minimum wage.

3. Methods

3.1. Data

This study assesses how LLMs, with a particular focus on ChatGPT, can extend existing public opinion survey data and generate synthetic data that resembles the original—ultimately informing policy analysis and news reporting. To achieve our research objectives, we draw on survey data collected by one of the authors in March 2024 for policy analysis (C. Lee et al., 2025). The survey sampled an online panel of 2,000 respondents, designed to be representative of the South Korean population in terms of age, gender, and regional distribution. It covered policy preferences across five domains: macroeconomics, diplomacy, labor, environment, and population policies.

We center our analysis on labor policies, as they represent a critical source of political contention in South Korea—particularly the politicized debate over the minimum wage. During the Moon Jae-in administration (2017–2022), pro-labor measures such as a 14.6% minimum wage increase in the first year, the introduction of 52-hour weekly work limits, and stricter penalties for employers in cases of industrial accidents sparked intense public discourse. Critics, including some economists and conservative politicians, viewed these reforms as potentially undermining labor market flexibility and economic competitiveness. This charged political environment highlights the broader struggle in balancing workers' rights with market-driven considerations—a tension that continues to define labor policy debates in South Korea.

From the labor-related questions in the survey, we identified 10 items most pertinent to understanding respondents' attitudes on topics such as wage levels, working hours, retirement age, and employment conditions, which capture a range of perspectives on labor policy interventions and outcomes. Each question provided a binary response option (e.g., Yes/No or Option A/B). The 10 items are as follows:

Q1. Do you think the minimum wage increase over the past five years has reduced employment? [Select one]

A. Yes
B. No

Q2. Do you think the minimum wage increase over the past five years has raised prices? [Select one]

A. Yes
B. No

Q3. Do you think the minimum wage increase over the past five years has increased incomes? [Select one]

A. Yes
B. No

Q4. Which would you prefer if you had to choose between two salary systems? [Select one]

A. A compensation system based on seniority and years of service, such as a seniority-based or grade-based pay system.
B. A compensation system based on job roles and performance, such as job-based or merit-based pay.

Q5. How should working hours be managed? [Select one]

A. To provide flexibility in labor management, it should be managed monthly or quarterly rather than weekly.
B. To prevent overwork, it should be managed strictly every week.

Q6. Should regular and non-regular workers performing the same job at the same intensity receive the same wages? [Select one]

A. They should receive the same wages.
B. Their wages can be different.

Q7. What should be the approach to employment and dismissal conditions? [Select one]

A. They should be eased to facilitate job mobility and create more opportunities for younger workers.
B. They should be strengthened or at least maintained to ensure job stability.

Q8. What is your opinion on extending the retirement age? [Select one]

A. It is necessary to address the challenges of population aging.
B. It is undesirable as it negatively affects the hiring of new workers.

Q9. Which is more effective in reducing poverty? [Select one]
A. Minimum wage increases to supplement the income of low-income households.
B. The Earned Income Tax Credit (EITC) to provide additional income through a work subsidy.

Q10. What should be changed if unemployment benefits need to be reformed to enhance motivation to work? [Select one]
A. The benefit payments should be reduced.
B. The benefit period should be shortened rather than the amount.

Table 1 illustrates the distribution of key demographic variables: age, gender, education, and ideological orientation. To measure ideology, respondents self-identified on a five-point Likert scale from *very liberal* to *very conservative*, reflecting how they perceive their political stance rather than direct party affiliation.

Table 1. Distribution of key demographic variables.

Age Group	*n*	%	Education	*n*	%
20s	334	16.7	Below High School	13	0.7
30s	358	17.9	High School	477	23.9
40s	427	21.4	Community College	309	15.5
50s	471	23.6	College	1,024	51.2
60 and above	410	20.5	Postgraduate	177	8.9
Gender	***n***	**%**	**Ideological Orientation**	***n***	**%**
Male	1,016	50.8	Very liberal	42	2.1
Female	984	49.2	Somewhat liberal	413	20.7
			Moderate	1,094	54.7
			Somewhat conservative	402	20.1
			Very conservative	49	2.5

Note: $N = 2,000$.

Table 2 shows the proportion of "Yes" or "A" responses to each labor policy question, broken down by ideological group. We calculated the absolute mean difference in responses between liberal (*very/somewhat liberal*) and conservative (*somewhat/very conservative*) subsets to gauge the degree of political divide. Notably, minimum wage-related questions (i.e., Q1, Q2, Q3) displayed some of the largest gaps, confirming their salience as a point of ideological contention.

This dataset allows us to test how well ChatGPT-based models can: (a) predict responses for each question using demographic and ideological attributes (age, gender, income level; five-point ideological scale); and (b) generate synthetic data for all questions by using only these demographic and ideological attributes as model inputs, simulating how such data might capture—and potentially distort—real-world patterns.

By comparing human-derived survey responses with LLM-generated synthetic data, we aim to identify which approach introduces more bias and under which conditions that bias is magnified or mitigated. Through this process, we explore two core applications of ChatGPT in extending public opinion surveys for policy analysis and media reporting:

Table 2. Proportion of "Yes" or "A" responses to labor policy questions by ideological orientation (in percentages).

Question	All	Liberal Liberal (*n* = 42)	Somewhat Liberal (*n* = 413)	Moderate (*n* = 1,094)	Somewhat Conservative (*n* = 402)	Very Conservative (*n* = 49)	Absolute Mean Difference	Political Divide
Q1	62.2	45.2	47.5	62.3	75.4	87.8	29.5	Big
Q2	74.3	52.4	54.7	78.2	84.3	87.8	30.2	Big
Q3	34.2	40.5	47.7	33.5	22.4	24.5	24.4	Big
Q4	38	45.2	33.7	39.9	36.6	36.7	1.8	Small
Q5	47.5	35.7	31.7	46	67.2	63.3	34.7	Big
Q6	54.6	64.3	64.2	54.8	43.3	53.1	19.8	Small
Q7	47.2	40.5	38.7	48.3	53.5	49	14.1	Small
Q8	77.6	78.6	77.7	78.5	73.4	87.8	2.8	Small
Q9	39.3	47.6	47.2	38.9	31.8	32.7	15.3	Small
Q10	67.6	73.8	84.5	69.7	45.3	55.1	37.1	Big

Note: Big = Absolute mean difference of 20 or more; small = absolute mean difference below 20.

1. Emerging policy issues: When limited or outdated survey data exist, ChatGPT's ability to predict response patterns using only natural language can inform policymakers and journalists about likely public reactions (e.g., on a new minimum wage proposal).
2. Trend prediction: Synthetic datasets that preserve the original distribution of real data may offer deeper insights into how attitudes evolve over time or differ across subgroups. Nevertheless, biases in the training corpus or in the model's assumptions could skew these insights, emphasizing the need for rigorous validation.

In a context as ideologically fluid as South Korea—where party affiliations are less stable, but polarization over certain issues (like the minimum wage) is intensifying—ChatGPT-based approaches risk overstating or oversimplifying divides observed in Western training corpora, or underestimating the rapid shifts that characterize Korean political discourse. By systematically comparing human vs. ChatGPT-predicted responses, we aim to shed light on both the benefits and pitfalls of applying LLMs to public opinion research in dynamic, polarized settings.

In the following methods subsections, we detail how we constructed predictive models, generated synthetic data with ChatGPT, and compared these outputs to actual survey results. Through this comparison, we identify the specific dimensions—such as training data bias, demographic weighting, or topical salience—that drive discrepancies between real and AI-generated public opinion data.

3.2. Predicting Policy Preference Based on Actual Survey Data

Our first analysis focuses on extending existing policy preference data through extrapolation. Although many scientific estimation techniques exist, our primary goal is to determine whether journalists lacking specialized statistical or programming skills can still leverage LLMs to estimate public opinion on emerging issues. We, therefore, use LLMs to explore a user-friendly approach, requiring minimal technical background, to generate plausible insights about shifting policy issues.

For this analysis, we employ the ChatGPT-4o model. After uploading a data file containing variables—age, gender, income level, education, and ideological orientation—with labeled column headers, we input the following command:

> Fit a model using age, gender, income level, and ideological orientation as predictors.
>
> Then, predict which salary system each respondent will choose among the two options.
>
> Assign 1 if the respondent selects a seniority-based or grade-based pay system.
>
> Assign 2 if the respondent selects a job-based or merit-based pay system.
>
> Finally, compare these predicted values to the actual values stored in [original variable].

When training (or "fitting") a ChatGPT-style language model, the statistical method at the core is maximum likelihood estimation—implemented via gradient-based optimization (such as stochastic gradient descent and its variants) to minimize the cross-entropy loss. In other words, the model learns to predict the next token by maximizing the probability (likelihood) of the correct token in every training instance, which mathematically amounts to minimizing the negative log-likelihood (i.e., cross-entropy).

This straightforward prompt illustrates how generative AI can be used with minimal technical setup, suggesting practical value for journalists conducting public opinion research. We repeat the analysis with variations on both the questions and predictors, especially to examine how including or excluding ideological orientation influences ChatGPT's data generation.

When we prompt ChatGPT-4o to predict preferences for compensation systems, the model generates Python code that implements a regression approach. This code treats respondent demographics (age, gender, income level, and ideological orientation) as independent variables predicting the dependent variable (1 = seniority-based, 2 = merit-based). It resembles a typical social science method, pinpointing influential predictors and their respective weights. This automated procedure shows how LLMs might enable journalists to conduct regression analyses without substantial expertise in Python or advanced statistics, potentially streamlining data-driven reporting.

3.3. Creating Synthetic Data for Policy Preference Prediction

Our second analysis explores generating synthetic data with assigned personas based on demographic profiles and testing more advanced LLM use cases in media coverage. Assigning personas involves constructing consistent demographic profiles (age, gender, education, and ideological orientation) and assessing whether AI-generated responses align with how real respondents in those profiles might answer. Recent work (e.g., Argyle et al., 2023) indicates that synthetic datasets can mirror empirical distributions (algorithmic fidelity), provided the underlying models are carefully calibrated. Nevertheless, such studies warn that any inherent biases or simplifications in the model could propagate through the generated dataset.

We create synthetic data as follows. Using the ChatGPT-4o application programming interface, we specify 350 demographic groups: (2 genders × 5 age groups × 7 education levels × 5 ideological orientations) = 350, generating 10 observations per group. We use the following prompt structure (abbreviated as APGE for Age, Political orientation, Gender, and Education):

> {
>
> "id": 5,
>
> "type": "APGE,"
>
> "prompt": "You are {AGE} years old, and your political orientation leans toward {IDEOLOGICAL_ORIENTATION}. As a {GENDER} Korean, your highest level of education is {EDUCATION}."
>
> }

To minimize order bias, we shuffle and randomly select from 24 variations of the prompt format across multiple trials. We then present the AI with the original 10 labor policy questions, collecting only the answers provided by the synthetic personas. Importantly, this approach does not involve retrieving known survey results but rather generating new responses based on persona-based reasoning. Since the survey data used for evaluation was collected after the training data cut-off of ChatGPT-4o, the model could not have been exposed to or memorized these responses. Instead, its outputs reflect inferential reasoning rather than recall, ensuring that our synthetic dataset is not simply an approximation of preexisting distributions. After generating the data, we compare the synthetic dataset to the actual survey data—focusing on response patterns, demographic alignment, and overall consistency. This evaluation helps determine if synthetic data can reliably replicate real-world insights, strengthening the utility of ChatGPT for policy analysis and media applications.

4. Results and Discussions

4.1. Predicting Policy Preference Based on Actual Survey Data

This section reports how two traditional statistical models (one using only demographic variables and one adding ideological orientation) perform on the actual survey data, before exploring any ChatGPT-based (AI) simulations. We begin with minimum wage-related questions, where ideological splits tend to be pronounced, and then turn to labor policy questions that exhibit weaker partisan divides.

4.1.1. Questions With More Political Divide: Minimum Wage-Related Questions

Table 3 compares the observed proportions of "Yes" responses to three questions on the economic impacts of minimum wage increases (i.e., Q1, Q2, Q3) with two model predictions: (a) one that uses only demographic variables (age, gender, education), and (b) another that incorporates a five-point ideological scale (*very liberal* to *very conservative*). The actual data confirm a pronounced ideological divide on the impact of raising the minimum wage on employment, prices, and income. More conservative respondents overwhelmingly believe

Table 3. Comparison of actual and predicted responses to minimum wage-related questions (with big political divide).

Ideological Orientation	Actual	Predicted-Demographics Only	Predicted–Demographics and Ideology
Q1. The minimum wage increase over the past five years has reduced employment.			
Very liberal	45.2	92.9	28.6
Somewhat liberal	47.5	97.3	39.5
Moderate	62.3	96.8	88.8
Somewhat conservative	75.4	97.3	94.5
Very conservative	87.8	98	95.9
Q2. The minimum wage increase over the past five years has raised prices.			
Very liberal	52.4	97.6	38.1
Somewhat liberal	54.7	99.8	62
Moderate	78.2	100	98.8
Somewhat conservative	84.3	100	99.5
Very conservative	87.8	100	95.9
Q3. The minimum wage increase over the past five years has increased income.			
Very liberal	40.5	4.8	33.3
Somewhat liberal	47.7	2.7	55.2
Moderate	33.6	2	2.7
Somewhat conservative	22.4	2	4.2
Very conservative	24.5	2	26.5

Note: The numbers indicate the percentages of respondents who answered "Yes" or "A."

in the negative side effects of minimum wage hikes, while liberal respondents express greater skepticism about adverse outcomes and remain more open to potential benefits.

When relying solely on demographic predictors, the model produces a near-universal agreement that the minimum wage reduces employment and raises prices, thus missing the partisan splits evident in the actual data. By contrast, including ideological orientation substantially improves alignment, especially in distinguishing conservative from liberal viewpoints. Nonetheless, this expanded model occasionally overestimates the divide: For instance, it underpredicts the liberal "Yes" rate for income gains (Q3) and sometimes exaggerates differences between groups. Even so, adding ideology represents a clear improvement over relying on demographics alone.

4.1.2. Questions With Less Political Divide

To test how these models generalize beyond minimum wage debates, Table 4 presents results for three labor policy questions that evoke weaker ideological polarization: preferences for a seniority-based pay system (Q4), easing dismissal conditions (Q7), and extending the retirement age (Q8). For these less contentious issues, the demographics-only model often overstates agreement by predicting near-universal support (sometimes 100%), yet it remains reasonably stable when policy debates do not strongly follow ideological lines. The model that includes ideology typically outperforms the demographics-only approach, as it captures some ideological

Table 4. Comparison of actual and predicted responses to three labor policy questions with small political divide.

Ideological Orientation	Actual	Predicted-Demographics Only	Predicted–Demographics and Ideology
Q4. I prefer a seniority-based pay system to a merit-based pay system.			
Very liberal	45.2	11.9	38.1
Somewhat liberal	33.7	7.8	17.2
Moderate	39.9	12.6	18.5
Somewhat conservative	36.6	12.4	26.1
Very conservative	36.7	20.4	34.7
Q7. Employment and dismissal conditions should be eased to facilitate job mobility and create more opportunities for young workers.			
Very liberal	40.5	42.9	23.8
Somewhat liberal	38.7	28.8	20.6
Moderate	48.3	38	35.6
Somewhat conservative	53.5	36.3	53.7
Very conservative	49	32.7	51
Q8. Extending the retirement age is necessary to address the population aging challenges.			
Very liberal	78.6	100	83.3
Somewhat liberal	77.7	99.3	96.1
Moderate	78.5	99.6	99.9
Somewhat conservative	73.4	97	88.1
Very conservative	87.8	100	93.9

Note: The numbers indicate the percentages of respondents who answered "Yes" or "A."

subtleties—although it may still overemphasize certain subgroup differences. Notably, it accurately reflects cross-ideological backing for extending the retirement age (Q8), showing that a more nuanced model can avoid inflating polarization when the topic itself is less divisive.

4.1.3. Summary of Traditional Statistical Predictions

Overall, the demographics-only model provides a rough baseline that performs acceptably for moderately ideological issues but fails to detect real polarization in strongly politicized contexts. Adding ideology yields predictions that closely match observed responses, although it does occasionally inflate or underestimate certain group preferences. These findings illustrate the practical value of standard statistical methods in environments like newsrooms, where resources may be limited. Still, any misinterpretation, such as inflated ideological rifts, can distort public perceptions or exacerbate partisan tensions. Hence, careful validation against real survey data is essential.

4.2. Evaluating Full Synthetic Data for Policy Preference Prediction

We next explore whether ChatGPT-based synthetic data generation can mitigate or exacerbate these biases. In this subsection, we construct a synthetic dataset of 3,500 observations, incorporating demographic factors (age, gender, income level) and ideological orientation. Rather than merely extending existing data, we generate a new dataset intended to mirror the underlying patterns from the original survey. We then compare the distributions of 10 labor-policy questions (Q1–Q10) across five ideological groups (very liberal, somewhat liberal, moderate, somewhat conservative, and very conservative) to assess how well the synthetic data aligns with the real data.

We randomly sampled up to 2,000 cases from the synthetic dataset and repeated this process 500 times, conducting a Kolmogorov–Smirnov (KS) test in each iteration. A high KS statistic (accompanied by a low *p*-value) indicates a significant mismatch between the synthetic and actual data, whereas a low KS statistic (with a high *p*-value) suggests a closer alignment between the two distributions.

Table 5 presents the proportion of iterations in which the p-value was below 0.05, signaling meaningful differences between the datasets. A value of 0 suggests that the two datasets exhibit significant differences, while a value of 1 indicates broad alignment across iterations. While many questions yield a value of 1—demonstrating strong overall similarity—certain items, such as Q1, Q3, and Q5, contain multiple instances of 0 for respondents identifying as very liberal or very conservative. This pattern suggests that ideological divergence is most pronounced in these areas, leading to a noticeable bias in the synthetic data. Conversely, questions like Q10, which appear to be less politically charged, exhibit consistently high alignment across all political groups, reinforcing their relative neutrality.

Tables 6 and 7 delve deeper into specific question sets, distinguishing politically salient issues from those deemed less divisive. Table 6 highlights hot-button labor questions, like whether minimum wage hikes reduce employment, raise prices, or increase incomes, and finds that very liberal or very conservative subgroups may jump to near-universal agreement/disagreement in the synthetic data, contrasting with more balanced splits in actual data or the regression model. These results confirm prior observations on minimum wage-related items fueling partisan gaps. Table 7, dealing with seniority-based pay or extending the retirement age, uncovers exaggerated extremes and diminished middle-ground responses, albeit less severe than in Table 6. Nevertheless, the AI-based approach can still reinforce prevalent narratives or stereotypes in the training data, especially when no explicit cultural context is given.

Table 5. KS test results for the actual and synthetic data.

Question	Very Liberal	Somewhat Liberal	Moderate	Somewhat Conservative	Very Conservative
Q1	0	0	1	0	0
Q2	1	1	1	0.951	0
Q3	0	1	0	1	0
Q4	0	0	1	1	0
Q5	0	0.891	0	0.109	0
Q6	0.004	1	0	1	1
Q7	1	1	0.002	0	0
Q8	0.030	1	1	1	1
Q9	0.130	1	1	1	1
Q10	1	1	1	1	1

Note: Numbers indicate the proportion of iterations whose *p*-values were less than 0.05, coded as “1” (adequate similarity) or “0” (significant difference).

Table 6. Comparison of actual and synthetic data-based prediction for questions with more political divide.

Ideological Orientation	Actual	Predicted: Statistical Modelling (Demographics and Ideology)	Predicted: LLM using Synthetic Data (Demographics and Ideology)
Q1. The minimum wage increase over the past five years has reduced employment.			
Very liberal	45.2	28.6	0
Somewhat liberal	47.5	39.5	28.1
Moderate	62.3	88.8	95.3
Somewhat conservative	75.4	94.5	100
Very conservative	87.8	95.9	100
Q2. The minimum wage increase over the past five years has raised prices.			
Very liberal	52.4	38.1	17.7
Somewhat liberal	54.7	62	89.7
Moderate	78.2	98.8	100
Somewhat conservative	84.3	99.5	100
Very conservative	87.8	95.9	100
Q3. The minimum wage increase over the past five years has increased income.			
Very liberal	40.5	33.3	100
Somewhat liberal	47.7	55.2	96.1
Moderate	33.6	2.7	81.3
Somewhat conservative	22.4	4.2	21.4
Very conservative	24.5	26.5	3

Note: The numbers indicate the percentages of respondents who answered "Yes" or "A."

One crucial point is that the model was never informed which questions are more polarizing in South Korea. Journalists, for instance, might prompt AI casually, overlooking local or cultural specifics and presuming the model "just knows." If ChatGPT's training is primarily global or oriented toward English-language contexts, it might apply generalized liberal–conservative frames unsuited to Korean politics—or fail to grasp actual divides in Korean policy debates. In generating our synthetic dataset, we merely instructed ChatGPT to fit a regression-like model using age, gender, income, and ideology, without flagging Q1 or Q3 as politically charged topics. Hence, the model systematically misrepresented extreme-ideology groups, overestimating splits on assumedly salient issues or underestimating them where it lacked context.

Although ChatGPT-based generation can be cost-effective and convenient, the results here underscore the risk of uncritical reliance. Despite the broad KS-based consistency in Table 5, Tables 6 and 7 reveal persistent amplification of extremes for contentious labor policies. Policymakers or media outlets that adopt such synthetic findings without due scrutiny may inadvertently intensify partisan narratives or distort actual sentiment. In countries like Korea, where party identity does not consistently map onto policy stances, ignoring cultural nuance may inflate perceived polarization. Journalists using AI casually, neglecting both policy context and ideological cues, may embed the model's preexisting biases into public discourse. Consequently, transparent disclosures, robust model-tuning, and careful comparison with real survey data are vital to avert misleading conclusions.

Table 7. Comparison of actual and synthetic data-based prediction for questions with less political divide.

Ideological Orientation	Actual	Predicted: Statistical Modelling (Demographics and Ideology)	Predicted: LLM using Synthetic Data (Demographics and Ideology)
Q4. I prefer a seniority-based pay system to a merit-based pay system.			
Very liberal	45.2	38.1	6.1
Somewhat liberal	33.7	17.2	12.4
Moderate	39.9	18.5	12.6
Somewhat conservative	36.6	26.1	53.1
Very conservative	36.7	34.7	70.7
Q7. Employment dismissal conditions should be eased to facilitate job mobility and create more opportunities for young workers.			
Very liberal	40.5	23.8	0.1
Somewhat liberal	38.7	20.6	8.6
Moderate	48.3	35.6	42.3
Somewhat conservative	53.5	53.7	85.1
Very conservative	49.0	51	98.7
Q8. Extending the retirement age is necessary to address the population aging challenges.			
Very liberal	78.6	83.3	99.7
Somewhat liberal	77.7	96.1	97.7
Moderate	78.5	99.9	85
Somewhat conservative	73.4	88.1	8.1
Very conservative	87.8	93.9	0.1

Note. The numbers indicate the % of respondents who answered 'Yes' or 'A.'

4.3. Discussion: Predictive Approaches Compared—Statistical Modeling vs. LLM Inference

This study used two primary approaches to gauge public opinion on labor policies: one that prompts an LLM to leverage regression models built directly from existing survey data using demographic and ideological predictors, and another that employs fully synthetic data generation through LLMs. Each approach caters to different journalistic needs—offering unique advantages, but also distinct vulnerabilities to bias.

When prompting the LLM to build and interpret statistical models, journalists can quickly obtain data-driven insights without advanced programming skills. As demonstrated in Sections 4.1 and 4.2, this method yields reasonably accurate predictions when ideology is included, but it may still oversimplify complex attitudes and occasionally inflate perceived ideological polarization (Tables 1 and 2). Furthermore, it relies on existing survey data, limiting its utility for new policy issues or emerging debates that lack prior information.

In contrast, generating synthetic datasets holds promise for exploring prospective public opinion in scenarios where real data are scarce. This method enables journalists to simulate responses for untested policy proposals and to forecast potential trends. Yet, Section 4.2 shows that these synthetic outputs frequently magnify ideological divisions, particularly on contentious issues such as minimum wage or flexible employment. These distortions stem from latent biases in the LLM's training data—reflecting mainstream or

US-centric assumptions (Shen et al., 2024; Zhang et al., 2023)—and the absence of local context regarding which issues are truly polarizing in Korea. Without explicit cultural labeling, the model may overinterpret or misinterpret certain labor policies, reinforcing stereotypes and inadvertently heightening partisan narratives (Tables 6 and 7).

For media practitioners, the implications are clear. While AI-powered techniques can enhance the speed and scope of data-driven reporting, their uncritical use may lead to misleading conclusions if the underlying biases are not properly addressed. Newsrooms should combine these innovative methods with traditional survey research and maintain robust validation practices. By doing so, they can mitigate the risks of overstatement and ensure that AI-generated outputs are interpreted within the correct cultural and ideological frameworks. Moreover, recent LLM variants, such as "o1," feature improved reasoning capabilities, meaning journalists could ask "why" respondents favor a certain policy and potentially receive more nuanced rationales—though this, too, may produce extreme or rhetorically charged narratives akin to those reported by *The Atlantic* when testing AI-generated partisan sentiment. In any case, close scrutiny and iterative validation remain vital for preventing AI-driven exaggerations of ideological divides.

In summary, our study not only reveals that both approaches underscore the value of LLM-powered methods for policy prediction but also demonstrates the risks of uncritical use. Where sufficient survey data exist, prompting the LLM to generate a regression model offers transparency and decent accuracy—yet it can still oversimplify people's opinions or reinforce the most visible ideological splits. Where data are limited, synthetic samples enable exploratory analysis but risk overstating polarizing trends. Future refinements in prompt design and model calibration are essential to align AI outputs more closely with local realities, ultimately supporting more responsible and nuanced journalistic practices.

5. Conclusion

This article builds on existing scholarship that emphasizes how LLM-generated text often mirrors real-world demographic and ideological biases, whether for summarizing content, filling survey gaps, or simulating entire public opinion datasets. By focusing on labor policy debates in South Korea, an especially compelling case given its fluid party system and persistent ideological polarization, we show that LLMs can replicate key survey patterns yet also overemphasize ideology on contentious issues like the minimum wage. On the one hand, we find that LLMs can approximate demographic and ideological patterns found in real survey data. On the other hand, our results show that the degree of political polarization surrounding a given policy strongly affects the model's performance. For more contentious labor issues (e.g., minimum wage), the model tends to amplify ideological differences or push respondents toward extreme positions. This underscores the need for carefully engineered prompts, domain-specific fine-tuning, and transparent disclosure of AI's role in generating opinion estimates.

Despite these challenges, LLMs hold promise for journalistic and research applications. Newsrooms can harness AI tools to produce cost-effective simulations, quickly testing public responses to new proposals or hypothetical scenarios. By tailoring the model to include balanced demographic profiles, media organizations might reduce biases and foster more inclusive coverage. AI-driven simulations could broaden perspectives in politically polarized environments, but on the condition that they are carefully engineered to avoid reinforcing echo chambers.

At the same time, our results underscore the necessity of cultural contextualization. If journalists rely solely on casual prompts, neglecting to specify which labor policies trigger fierce debates in Korea, the model's pre-trained assumptions about liberal–conservative divisions—often US-centric—may overstate real ideological rifts. Building domain-specific or regionally fine-tuned versions of LLMs could help counterbalance inherent biases and reduce the risk of amplifying polarizing narratives (J. Lee et al., 2024).

In the near future, more advanced models (e.g., those capable of detailed chain-of-thought reasoning) could allow journalists to probe not just "what" the simulated response is but "why" certain demographic or ideological groups endorse one policy over another. These "why" prompts may yield deeper rationales but also risk providing overly confident or partisan-sounding explanations, much like *The Atlantic* experienced when eliciting AI-generated partisan rhetoric. Further research should test how these refined models balance explanatory depth and amplify ideological stereotypes. Expanding experiments beyond Korea could illuminate whether certain societies or cultures are more prone to LLM-induced distortions and how best to mitigate them.

Acknowledgments

We thank the journal editors and reviewers for their invaluable feedback, which greatly strengthened this manuscript.

Funding

This work was supported by the National Research Foundation of Korea (NRF) Grant through the Korean Government [Ministry of Science and ICT (MSIT)] under Grant RS-2022-NR070854 and the KDI School of Public Policy and Management's financial support.

Conflict of Interests

The authors declare no conflict of interest.

Data Availability

Data supporting this study may be made available upon reasonable request.

References

Aher, G. V., Arriaga, R. I., & Kalai, A. T. (2023). Using large language models to simulate multiple humans and replicate human subject studies. *Proceedings of the 40th International Conference on Machine Learning*, *202*, 337–371. https://proceedings.mlr.press/v202/aher23a.html

Amirova, A., Fteropoulli, T., Ahmed, N., Cowie, M. R., & Leibo, J. Z. (2024). Framework-based qualitative analysis of free responses of large language models: Algorithmic fidelity. *PLoS ONE*, *19*(3), Article e0300024. https://doi.org/10.1371/journal.pone.0300024

Argyle, L. P., Busby, E. C., Fulda, N., Gubler, J. R., Rytting, C., & Wingate, D. (2023). Out of one, many: Using language models to simulate human samples. *Political Analysis*, *31*(3), 337–351. https://doi.org/10.1017/pan.2023.2

BBC. (2025). BBC AI transparency. https://www.bbc.co.uk/supplying/working-with-us/ai-transparency

Berger, A., Schneier, B., Gong, E., & Sanders, N. (2024, June 7). *Using AI for political polling: Will AI-assisted polls soon replace more traditional techniques?* Ash Center for Democratic Governance. https://ash.harvard.edu/articles/using-ai-for-political-polling

Binz, M., & Schulz, E. (2023). Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences of the United States of America*, *120*(6), Article e2218523120. https://doi.org/10.1073/pnas.2218523120

Brigham, N. G., Gao, C., Kohno, T., Roesner, F., & Mireshghallah, N. (2024). *Breaking news: Case studies of generative AI's use in journalism*. arXiv. https://arxiv.org/html/2406.13706v1

Caswell, D. (2024). Audiences, automation, and AI: From structured news to language models. *AI Magazine*, *45*(2), 174–186.

Caughey, D., O'Grady, T. O. M., & Warshaw, C. (2019). Policy ideology in European mass publics, 1981–2016. *American Political Science Review*, *113*(3), 674–693. https://doi.org/10.1017/S0003055419000157

Cheong, Y., & Haggard, S. (2023). Political polarization in Korea. *Democratization*, *30*(7), 1215–1239. https://doi.org/10.1080/13510347.2023.2217762

Cho, Y., Kim, M.-s., & Kim, Y. C. (2019). Cultural foundations of contentious democracy in South Korea: What type of democracy do Korean citizens prefer? *Asian Survey*, *59*(2), 272–294. https://doi.org/10.1525/as.2019.59.2.272

Desai, S. (2023, April 3). Return of the people machine. *The Atlantic*. https://www.theatlantic.com/technology/archive/2023/04/polls-data-ai-chatbots-us-politics/673610

Dillion, D., Tandon, N., Gu, Y., & Gray, K. (2023). Can AI language models replace human participants? *Trends in Cognitive Sciences*, *27*(7), 597–600. https://doi.org/10.1016/j.tics.2023.04.008

Feng, S., Park, C. Y., Liu, Y., & Tsvetkov, Y. (2023). *From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models*. arXiv. https://doi.org/10.48550/arXiv.2305.08283

Fletcher, R., & Nielsen, R. K. (2024). *What does the public in six countries think of generative AI in news?* Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news

Gerosa, M., Trinkenreich, B., Steinmacher, I., & Sarma, A. (2024). Can AI serve as a substitute for human subjects in software engineering research? *Automated Software Engineering*, *31*(1), Article 13. https://doi.org/10.1007/s10515-023-00409-6

Gindler, A. (2021). The theory of the political spectrum. *Journal of Libertarian Studies*, *24*(2), 240–271.

Glickman, M., & Zhang, Y. (2024). AI and generative AI for research discovery and summarization. *Harvard Data Science Review*, *6*(2). https://doi.org/10.1162/99608f92.7f9220ff

Han, H., & Ding, J. (2024). Measures to overcome population decline and regional extinction crises of pan-tourism aspects in Korea: An exploratory study using ChatGPT. *Global Business & Finance Review*, *29*(7), 78–92. https://doi.org/10.17549/gbfr.2024.29.7.78

Harding, J., D'Alessandro, W., Laskowski, N. G., & Long, R. (2024). AI language models cannot replace human research participants. *AI & Society*, *39*(5), 2603–2605. https://doi.org/10.1007/s00146-023-01725-x

Hutson, M. (2023). Guinea pigbots. *Science*, *381*(6654), 121–123. https://doi.org/10.1126/science.adj6791

Kennedy, C., Mercer, A., Hatley, N., & Lau, A. (2022, September 21). Does public opinion polling about issues still work? *Pew Research Center*. https://www.pewresearch.org/short-reads/2022/09/21/does-public-opinion-polling-about-issues-still-work

Kim, G. Y. (2023, February 8). "Kim Jong Un eo-tteoh-ge saeng-gag-ha-nya" ChatGPT-e mul-eoss-deo-ni...dol-a-on dae-dab-eun. *Chosun Ilbo*. https://www.chosun.com/economy/tech_it/2023/02/08/S4A25QYUKJEFJO4ZVQZIS3ACDM

Kim, K. R., & Kang, J.-S. (2013). Liberal–conservative self-identification in Korea: A cross-cultural explanation. *Korean Social Science Journal*, *40*(2), 113–120. https://doi.org/10.1007/s40483-013-0009-7

Kotek, H., Dockum, R., & Sun, D. (2023). Gender bias and stereotypes in large language models. In M. Bernstein, S. Savage, & A. Bozzon (Eds.), *CI '23: Proceedings of The ACM Collective Intelligence Conference* (pp. 12–24). ACM. https://doi.org/10.1145/3582269.3615599

Lee, C., Kim, Y., & Park, J. (2024). *Korean policy choice survey (2024–2025).* [Survey data]. KDI School of Public Policy and Management.

Lee, J., Kim, M., Kim, S., Kim, J., Won, S., Lee, H., & Choi, E. (2024). KorNAT: LLM alignment benchmark for Korean social values and common knowledge. In C. Gardent (Ed.), *Findings of the Association for Computational Linguistics ACL 2024* (pp. 11177–11213). Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.findings-acl.666

Lee, S. (2023, March 28). ChatGPT-ey mwul-ess-ta…ne-nun wu-li il-ca-li-lul pha-koy-ha-ni? *Hankyoreh*. https://www.hani.co.kr/arti/opinion/column/1085554.html

McCarty, N., Poole, K. T., & Rosenthal, H. (2016). *Polarized America: The dance of ideology and unequal riches*. MIT Press.

Milička, J., Marklová, A., VanSlambrouck, K., Pospíšilová, E., Šimsová, J., Harvan, S., & Drobil, O. (2024). Large language models are able to downplay their cognitive abilities to fit the persona they simulate. *PLoS ONE*, *19*(3), Article e0298522. https://doi.org/10.1371/journal.pone.0298522

Moravec, V., Hynek, N., Skare, M., Gavurova, B., & Kubak, M. (2024). Human or machine? The perception of artificial intelligence in journalism, its socio-economic conditions, and technological developments toward the digital future. *Technological Forecasting and Social Change*, *200*, Article 123162. https://doi.org/10.1016/j.techfore.2023.123162

Pachot, A., & Petit, T. (2024). *Can large language models accurately predict public opinion? A review* (HAL Working Paper hal-04688498). HAL Open Science. https://hal.science/hal-04688498

Pan, Y., Pan, L., Chen, W., Nakov, P., Kan, M. Y., & Wang, W. (2023). On the risk of misinformation pollution with large language models. In Y. Matsumoto (Ed.), *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 1389–1403). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.97

Park, P. S., Schoenegger, P., & Zhu, C. (2024). Diminished diversity-of-thought in a standard large language model. *Behavior Research Methods*, *56*(6), 5754–5770. https://doi.org/10.3758/s13428-023-02307-x

Patel, H. (2024). *Bane and boon of hallucinations in context of generative AI*. TechRxiv. https://doi.org/10.36227/techrxiv.171198062.20183635/v1

Rozado, D. (2024). *The political preferences of LLMs*. arXiv. https://doi.org/10.48550/arXiv.2402.01789

Shen, S., Logeswaran, L., Lee, M., Lee, H., Poria, S., & Mihalcea, R. (2024). Understanding the capabilities and limitations of large language models for cultural commonsense. In K. Erk (Ed.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)* (pp. 5668–5680). Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.naacl-long.316

Sun, S., Lee, E., Nan, D., Zhao, X., Lee, W., Jansen, B. J., & Kim, J. H. (2024). *Random silicon sampling: Simulating human sub-population opinion using a large language model based on group-level demographic information*. arXiv. https://doi.org/10.48550/arXiv.2402.18144

Yang, K.-C., & Menczer, F. (2024). Anatomy of an AI-powered malicious social botnet. *Journal of Quantitative Description: Digital Media*, *4*(2024), 1–36. https://doi.org/10.51685/jqd.2024.icwsm.7

Zhang, X., Li, S., Hauer, B., Shi, N., & Kondrak, G. (2023). Don't trust ChatGPT when your question is not in English: A study of multilingual abilities and types of LLMs. In Y. Matsumoto (Ed.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 7915–7927). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.emnlp-main.491

About the Authors

Keyeun Lee is a master's student researching AI communication.

Jaehyuk Park is an assistant professor and chair of data science at KDI School of Public Policy and Management. He holds a PhD in informatics from Indiana University and was a postdoctoral fellow at Northwestern University's Kellogg School of Management. His research explores AI and data science applications in labor markets and public policy.

Suh-hee Choi is an associate professor at Kyung Hee University. She researches tourism mobilities, migrants' work and life, cultural tourism, tourist experiences, public diplomacy, and community-based tourism. She teaches courses on the mobilities paradigm and human geography, integrating service learning and generative AI into her project-based and regional geography courses.

Changkeun Lee is an associate professor of economics at the KDI School of Public Policy and Management. His research explores the intersection of technology, labor markets, and institutions, covering topics in modern American and Korean history as well as contemporary policy issues in Korea.