

Will AI Take My Job? Evolving Perceptions of Automation and Labor Risk in Latin America

Andrea Cremaschi, Dae-Jin Lee, Manuele Leonelli

School of Science and Technology, IE University, Madrid, Spain
andrea.cremaschi@ie.edu, dae-jin.lee@ie.edu, manuele.leonelli@ie.edu

Abstract

As artificial intelligence and robotics increasingly reshape the global labor market, understanding public perceptions of these technologies becomes critical. We examine how these perceptions have evolved across Latin America, using survey data from the 2017, 2018, 2020, and 2023 waves of the Latinobarómetro. Drawing on responses from over 48,000 individuals across 16 countries, we analyze fear of job loss due to artificial intelligence and robotics. Using statistical modeling and latent class analysis, we identify key structural and ideological predictors of concern, with education level and political orientation emerging as the most consistent drivers. Our findings reveal substantial temporal and cross-country variation, with a notable peak in fear during 2018 and distinct attitudinal profiles emerging from latent segmentation. These results offer new insights into the social and structural dimensions of AI anxiety in emerging economies and contribute to a broader understanding of public attitudes toward automation beyond the Global North.

Extended version — <https://arxiv.org/abs/2505.08841>

Introduction

As artificial intelligence (AI) technologies become increasingly embedded in the fabric of economic and social life, public concern about their long-term implications has grown. Among the most salient fears is automation displacing human labor, particularly in economies marked by informality, inequality, and institutional fragility. While techno-optimist narratives highlight gains in productivity and efficiency, they often overlook the uneven social consequences of AI deployment, especially fears about job loss, economic marginalization, and weakened democratic accountability.

Existing research has shown that public attitudes toward AI are shaped by a complex interplay of individual characteristics, political orientations, and broader cultural and institutional contexts (Bergdahl et al. 2023; Gillespie, Lockey, and Curtis 2021; Schiavo, Businaro, and Zancanaro 2024). Several studies have found that education, age, gender, and political ideology influence perceptions of AI (Gerlich 2023; Zhang and Dafoe 2019; Kaya et al. 2024), while others emphasize the role of institutional trust in shaping public accep-

tance (Gillespie, Lockey, and Curtis 2021; Schiavo, Businaro, and Zancanaro 2024). However, most of this evidence is concentrated in North America, Europe, or East Asia. Relatively little is known about how populations in Latin America, where economic precarity and political polarization are prevalent, perceive the labor risks associated with AI.

This paper addresses this gap by analyzing responses from over 48,000 individuals across 16 Latin American countries, drawn from the four most recent waves (2017, 2018, 2020, and 2023) of the Latinobarómetro survey (Latinobarómetro Corporation 2017–2023). Using a combination of statistical modeling and latent class analysis, we examine how fear of AI-induced job loss varies across countries, time periods, and social groups. Our analysis highlights education level and political ideology as the most consistent predictors of concern, while also uncovering how these attitudes cluster into distinct societal profiles.

We find that concern about AI-driven job loss peaked in 2018, likely due to changes in question framing and heightened media attention, but remains widespread and unequally distributed. Fear is especially pronounced among individuals with lower educational attainment. Political ideology also plays a role, with left-leaning respondents expressing greater concern in earlier years, though this pattern becomes less consistent over time and, in some cases, reverses. The results underscore the importance of structural conditions and ideological identity in shaping public perceptions of automation in Latin America. In doing so, this study contributes a comparative perspective to a literature still heavily focused on the Global North, offering new insight into how technological anxiety is patterned in emerging economies.

Literature Review

Public concern about AI displacing human workers is a dominant theme in public discourse, policy debates, and empirical studies. However, how this fear varies across countries and social groups, and how it relates to institutional trust and economic perceptions, remains an active area of research. In this section, we review the literature across five main themes relevant to our study: general fears of job loss due to AI, geographic differences, sociodemographic variation, institutional trust, and economic outlook.

Fear of Job Loss to AI

Fears about AI-induced job loss are widespread. Frey and Osborne (2017) estimated that up to 47% of U.S. jobs could be at risk of automation, while Acemoglu and Restrepo (2020) estimated annual job losses of 360,000–670,000 due to robotics. Recent work confirms these concerns are now deeply embedded in public consciousness. Gerlich (2023) and Kaya et al. (2024) report that the threat of job displacement ranks among the most salient AI-related fears (Chiarini et al. 2023; Li and Huang 2020; Nnamdi, Ogunlade, and Abegunde 2023; Soueidan and Shoghari 2024; Wilde 2025). In the 2024 UK Tracker Survey, over 70% of respondents voiced concern about AI-induced unemployment (Department for Science, Innovation and Technology 2024). However, Schiavo, Businaro, and Zancanaro (2024) show that job replacement anxiety, while common, is not a strong predictor of AI acceptance (Uçar, Çapuk, and Yiğit 2025).

Geographic and Cultural Variation

The intensity of AI-related fears varies across regions. European respondents tend to be more skeptical of AI than their counterparts in emerging economies (Gillespie, Lockey, and Curtis 2021). Bergdahl et al. (2023) report that AI positivity is highest in Finland and Poland and lowest in France, suggesting clear cultural and institutional differences. Latin America remains understudied, but evidence from the *Latino-barómetro* surveys reveals a distinct pattern: concern about AI-driven job loss peaked sharply in 2018 before stabilizing, as we demonstrate in our longitudinal analysis below. This trajectory mirrors findings from the UK and Germany, where concern tends to rise in response to media and policy events but attenuates with increased familiarity and exposure (Dong et al. 2024; Gerlich 2023).

Sociodemographic Predictors of AI Fear

Age. Studies find that younger individuals hold more positive attitudes toward AI (Zhang and Dafoe 2019; Bergdahl et al. 2023), yet this effect may be mediated by digital literacy and perceived self-efficacy rather than age itself (Schiavo, Businaro, and Zancanaro 2024). However, some studies suggest younger individuals may also perceive greater long-term vulnerability to automation (Morikawa 2017).

Gender. Women consistently express greater anxiety about AI, particularly regarding fairness, job loss, and ethical misuse (Bergdahl et al. 2023; Kaya et al. 2024). Kaya et al. (2024) show that women report significantly higher levels of AI learning and job replacement anxiety.

Education. Education is a robust predictor of AI attitudes. Higher education is associated with lower anxiety and greater acceptance (Schiavo, Businaro, and Zancanaro 2024). AI literacy, defined as understanding, using, and critically evaluating AI, is both shaped by and mediates the effect of education on AI acceptance (Lund, Mannuru, and Agbaji 2024).

Ideology. Political ideology shapes concerns about AI. Right-leaning individuals tend to focus on cultural threats and loss of control, whereas left-leaning respondents emphasize labor rights and inequality (Gerlich 2023). These find-

ings echo broader ideological divides in technology governance (Bergdahl et al. 2023).

City Size and Technological Exposure. Existing research suggests that residents of larger cities are generally more exposed to AI and tend to hold more nuanced views, often combining optimism with concern (Department for Science, Innovation and Technology 2024). In contrast, individuals living in smaller towns or rural areas frequently express greater skepticism, likely due to lower digital engagement and limited perceived benefits.

Trust in Institutions and Governance

Trust in those developing and governing AI is a strong predictor of public acceptance. While trust in academia and healthcare is relatively high, governments and tech companies are viewed more skeptically (Gillespie, Lockey, and Curtis 2021). This trust gap fuels anxiety: respondents want regulation but question government competence (Cremaschi, Lee, and Leonelli 2025; Gerlich 2023; Montag et al. 2023), and anxiety drops when governance is trusted (Schiavo, Businaro, and Zancanaro 2024; Zhan et al. 2024). A lack of trust heightens concerns about autonomy, transparency, and job loss.

Economic Outlook and Risk Perceptions

Economic perceptions, both macro and personal, play a significant role in AI fear. Individuals who believe the economy is deteriorating are more likely to see AI as a threat to job security (Kaya et al. 2024). Bergdahl et al. (2023) report that negative income expectations are strongly correlated with AI negativity. However, economic optimism can mitigate AI fear. For example, individuals who see AI as economically beneficial are more likely to accept it even when aware of its risks (Gerlich 2023; George 2024). Thus, economic context mediates whether automation is viewed as a threat or a tool (Gu et al. 2025).

In summary, the literature shows that fear of AI-driven job loss is widespread but shaped by demographic, cultural, and institutional factors. Education and trust, especially in government, emerge as key buffers against AI anxiety, while negative economic expectations heighten it. These insights frame our investigation of job loss fears in Latin America, where structural inequality and institutional fragility may shape attitudes in distinct ways.

Materials and Methods

Dataset Description

The analysis is based on survey data collected by *Latino-barómetro*, a public opinion study conducted across Latin America. We use the four most recent editions of the survey (2017, 2018, 2020, and 2023), which include responses from 16 countries: Bolivia, Brazil, Chile, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Honduras, Mexico, Panama, Paraguay, Peru, Uruguay, and Venezuela. Each wave captures perceptions of democracy, governance, economic outlook, and, crucially, emerging technologies like AI.

Variable	Description	Levels
fear_ai	Fear of losing one’s job due to AI/robots	<i>Strongly disagree, Disagree, Agree, Strongly agree</i>
country	Country of residence	<i>Bolivia, Brazil, ..., Venezuela</i>
age	Age of respondent	Numeric (continuous)
gender	Gender	<i>Female, Male</i>
education	Education level	<i>Low, Medium, High</i>
city_size	Size of locality	<i>Small town, Mid-sized city, Large city, Capital</i>
econ_now	Perception of current economy	<i>Bad, Average, Good</i>
econ_future	National economic outlook	<i>Worse, Same, Better</i>
econ_personal_future	Personal/family economic outlook	<i>Worse, Same, Better</i>
trust_people	Generalized trust	<i>Cannot be too careful, Trust most</i>
democracy_pref	Support for democracy	<i>Doesn’t matter, Authoritarian OK, Democracy best</i>
demo_satisfaction	Satisfaction with democracy	<i>Not at all satisfied, Not very satisfied, Satisfied</i>
trust_parliament	Trust in parliament	<i>None, Little, Some or A lot</i>
trust_gov	Trust in government	<i>None, Little, Some or A lot</i>
trust_judiciary	Trust in judiciary	<i>None, Little, Some or A lot</i>
ideology_group	Political ideology	<i>Left, Centre, Right</i>
year	Year of response	<i>2017, 2018, 2020, 2023</i>

Table 1: Variables used in the analysis, with descriptions and categorical levels after final cleaning.

We analyze a subset of variables that are harmonized across waves and relevant to perceptions of AI. The outcome of interest is fear of job loss due to AI or robots, measured on a four-point Likert scale ranging from “Strongly disagree” to “Strongly agree.” Key predictors include demographic characteristics (age, gender, education, and city size), political orientation (left–centre–right self-placement), economic perceptions (views of the national economy, future national outlook, and future personal situation), and institutional attitudes (trust in people, government, parliament, and judiciary, as well as preferences for democracy and satisfaction with its functioning), with country and year included to capture geographic and temporal variation. Table 1 summarizes the chosen variables.

The question used to construct **fear_ai** varies slightly across waves. In 2017, respondents were asked whether artificial intelligence and robots would reduce more jobs than they create. In 2018, they were asked whether robots would take their job or a family member’s job. In 2020, the question was whether robots would take away their workplace within 10 years, and in 2023, whether robots would take their job in 10 years’ time. While the exact wording differs across waves, all questions tap into a shared concern about the future impact of robotics and AI on employment. We treat these questions as conceptually equivalent for the purpose of harmonized modeling, while acknowledging the slight variation in framing.

Data Preparation

We apply a standardized preprocessing pipeline to ensure comparability across survey waves, including harmonizing variable names and formats, recoding categorical responses into consistent factor levels, merging sparse or semantically similar categories, and enforcing ordinal structures where applicable. All variables except age are encoded as factors using sum-to-zero contrast coding, and we conduct a complete-case analysis to retain only fully observed records. Two countries are excluded: Nicaragua, not included in the

2023 wave, and Argentina, due to systematic missingness in key variables such as education and ideology in 2020. The final dataset comprises 48,403 individuals.

Exploratory Analysis

We conduct all statistical analyses using R within the RStudio environment (Posit team 2024). We first perform exploratory data analysis to assess the distribution and bivariate associations of the response variable **fear_ai** with the available predictors. We use the `ggplot2` package (Wickham 2016) to produce stacked barplots displaying the distribution of **fear_ai** across countries and years. Visualizations are faceted by country to aid comparison.

To quantify the strength of association between **fear_ai** and each feature, we compute mutual information (MI) using the `infotheo` package (Meyer 2022). To evaluate statistical dependence, we conduct Fisher’s exact tests of independence with simulated p-values. For the continuous predictor **age**, we discretize it into quartiles.

Statistical Modeling

To model the relationship between the outcome variable **fear_ai** and the available predictors, we use *ordinal logistic regression*, also known as the *proportional odds model* (Agresti 2013; McCullagh 1980). This approach is well suited for outcomes measured on a four-point Likert scale.

In ordinal logistic regression, the model estimates the log-odds of being at or below a certain category, given a set of predictor variables. Formally, let Y_i denote the ordinal response for individual i , taking values in $\{1, 2, \dots, K\}$, with $K = 4$ here. The model assumes:

$$\log \left(\frac{\mathbb{P}(Y_i \leq k)}{\mathbb{P}(Y_i > k)} \right) = \theta_k - \mathbf{x}_i^\top \boldsymbol{\beta}, \quad \text{for } k = 1, 2, \dots, K - 1,$$

where θ_k are $K - 1$ threshold parameters separating the response categories, \mathbf{x}_i is the vector of predictors for individual i , and $\boldsymbol{\beta}$ is the vector of regression coefficients. The

model assumes that β is constant across the $K - 1$ equations: this is known as the *proportional odds assumption*. We implement this model using the `polr()` function from the MASS package in R (Venables and Ripley 2002).

All categorical predictors are encoded using *sum-to-zero* constraints. This coding ensures that for any categorical variable with J levels, the model includes J coefficients whose sum equals zero. Each level's coefficient represents a deviation from the average effect of that variable, making comparisons across levels and countries more interpretable. Coefficient signs indicate direction: positive values imply greater likelihood of stronger fear responses, while negative values suggest the opposite. The magnitude of the coefficient reflects the strength of this relationship.

We fit separate models for each country and each year (2017, 2018, 2020, 2023), as well as a pooled "FULL" model per year that includes all countries. In each case, we use a *greedy forward selection* algorithm to identify a parsimonious set of predictors. Starting from an intercept-only model, we iteratively add the variable that yields the largest improvement in out-of-sample log-likelihood, measured using 10-fold cross-validation. This procedure continues until no remaining variable improves model fit.

To quantify the strength of association between a predictor and the outcome, we use the *average absolute coefficient* across all levels of the predictor. This measure provides an intuitive summary of a variable's impact: the larger the average absolute coefficient, the greater the predictor's contribution to shifting individuals across fear categories.

This framework allows us to examine both whether a variable is selected for inclusion in a given country-year model, and how strongly it influences attitudes toward AI-induced job loss.

Latent Class Analysis for Profile Identification

To uncover distinct latent profiles based on respondents' characteristics and attitudes, we perform Latent Class Analysis (LCA). LCA is a probabilistic clustering method that identifies subgroups of individuals who share similar response patterns across categorical variables, under the assumption of conditional independence within classes (Lazarsfeld 1950). We include all predictors used in the study, except for **country** and **fear.ai**, so that the latent classes reflect general patterns of demographic background, economic perceptions, political attitudes, and institutional trust, rather than being directly driven by geography or fear of AI-related job loss. The continuous variable **age** is again divided into quartiles.

LCA models the joint distribution of observed categorical variables $\mathbf{X} = (X_1, X_2, \dots, X_J)$ through a categorical latent variable $C \in \{1, \dots, K\}$, where K is the number of latent classes. Under the assumption of conditional independence, the model is specified as:

$$\mathbb{P}(\mathbf{X} = \mathbf{x}) = \sum_{k=1}^K \pi_k \prod_{j=1}^J \mathbb{P}(X_j = x_j \mid C = k),$$

where $\pi_k = \mathbb{P}(C = k)$ is the proportion of individuals in class k , and $\mathbb{P}(X_j = x_j \mid C = k)$ is the item-response prob-

ability: the probability of a particular response on variable X_j , conditional on membership in class k .

These probabilities are estimated via maximum likelihood using the `poLCA` package in R (Linzer and Lewis 2011). For each number of latent classes from 1 to 10, we fit models using 20 random starts to ensure robust convergence. For each model, we compute the Bayesian Information Criterion (BIC), which balances model fit and complexity, and the entropy, which quantifies classification certainty, with higher entropy indicating clearer separation between classes. We use both metrics to select a number of classes that fits the data well while producing distinct and interpretable profiles.

After selecting the optimal model, respondents are assigned to the class with the highest posterior membership probability. The estimated item-response probabilities are then used to characterize the substantive meaning of each latent profile. These classes form the basis for analyzing variation in AI-related job loss fear across societal segments. To assess whether **fear.ai** distributions differ significantly between classes, we conduct chi-squared tests for each survey year, followed by pairwise comparisons with adjustment for multiple testing (Benjamini and Hochberg 1995).

Results

Overview of Fear of AI-Driven Job Loss

Across the four waves of the Latinobarómetro survey, respondents expressed varying degrees of concern about the risk of job loss due to AI. Overall, the most common response was *Agree* (45.6%), followed by *Disagree* (29.6%). A smaller proportion of respondents reported *Strongly agree* (17.5%), while only 7.3% selected *Strongly disagree*, suggesting that fears related to AI are widespread but often moderate in intensity.

When examining responses by year, there is notable variation. In 2018, concern was especially pronounced: nearly 69.3% of respondents *Agreed* that robots could take their job or a family member's, while only 14.3% *Strongly agreed* and no respondents selected *Strongly disagree*. By 2020 and 2023, a more balanced spread across categories is observed, with both waves showing over 12% of respondents selecting *Strongly disagree*, and moderate concern (*Disagree* or *Agree*) distributed more evenly. Interestingly, in 2017, there was a greater proportion of respondents who expressed *Strongly agree* (28.0%) compared to any other year, pointing to an earlier polarization in views around AI-driven job displacement.

The evolution of concern about AI-driven job loss shows clear cross-country differences and dynamic patterns over time, as illustrated in Figure 1. Brazil stands out for its high polarization: across all waves, it exhibits a notably large share of both *Strongly agree* and *Strongly disagree* responses, much more so than any other country. This points to entrenched divisions over automation in Brazil's labor market. In contrast, countries such as Venezuela and Paraguay show initially high levels of fear, particularly in 2017 and 2018, followed by a noticeable decline in 2020. This pattern may reflect shifting perceptions as automation discourse matured or became less salient during that period. Most other




Figure 1: Fear of AI-driven job loss across years, by country. Bars show the distribution of responses to **fear_ai** from 2017 to 2023.

countries (with the notable exception of Brazil) display a U-shaped trend: relatively high concern in 2018, a dip in 2020, and a modest rebound in 2023. This dip may be partially attributed to a change in question framing. In 2018, respondents were asked whether robots could take their job *or a family member's job*, potentially inflating perceived risk, whereas in 2020 and 2023, the question focused exclusively on personal job loss. Nonetheless, the renewed rise in concern by 2023 suggests that fear of automation is resurging, likely fueled by increasing real-world adoption of AI technologies and heightened media attention.

Overall, while temporal patterns are broadly consistent across countries, the magnitude of fear and the degree of polarization vary substantially, underscoring the importance of contextual factors in shaping public perceptions of AI-driven labor market disruptions.

Associations Between Predictors and Fear of AI

Table 2 reports the association between **fear_ai** and each predictor, as measured by mutual information (MI) and tested via Fisher's exact test. The features most strongly associated with concern about AI-driven job loss are **education** and **country**, followed by **econ_future** and **demo_satisfaction**. These variables exhibit consistently higher MI values and highly significant p-values across all years.

Discretized **age_group** shows modest association with **fear_ai**, particularly in 2018, but is weaker in other years. In contrast, variables like **gender** and **trust_people** exhibit very low MI values despite occasionally significant p-values, suggesting that these variables contribute little explanatory power, despite some statistical significance. Overall, the results highlight that structural (e.g., education, economic outlook) and contextual (e.g., country, political ideology) factors are more informative than basic demographics or general attitudes.

Analysis of Covariate Effects Across Years

Figure 2 presents heatmaps summarizing the average absolute effect sizes of selected predictors across countries and years, based on cross-validated ordinal regression models. Darker shading corresponds to stronger covariate effects. There is a clear temporal pattern: in 2017, many covariates exhibited non-negligible effects across multiple countries, suggesting a broader range of predictors associated with AI-related job loss fear. In later years (2018, 2020, and 2023), the heatmaps are sparser with fewer variables consistently selected and smaller average effect sizes overall. This suggests that over time, attitudes toward AI job displacement became structured around a smaller set of key predictors.

Across all years, **education** emerges as the most robust and consistently strong predictor, appearing with relatively high intensity in nearly all countries and years. **Ideology_group** also maintains substantial importance across time, indicating that political orientation plays a durable role in shaping views about AI and automation.

Gender shows moderate importance in 2017 but diminishes almost entirely in later years, suggesting that gender differences in AI fear diminished over time. Conversely, **city_size** displays relatively stable importance across all four waves, particularly in countries such as Uruguay, Costa Rica, and Mexico, where urban-rural divides appear to influence perceptions of technological disruption.

Interestingly, **democracy_pref** has strong effects in 2017 but almost disappears as a predictor in subsequent years. This pattern may reflect changing political narratives around technology, or the relative decline in the salience of democratic preferences for structuring attitudes toward AI after the initial surge of attention in 2017–2018.

Comparing across countries, certain nations exhibit distinct profiles. For instance, Chile consistently shows strong effects for trust-related variables, including **trust_gov** and

Feature	Overall	2017	2018	2020	2023
country	0.0112 (0.000)	0.0321 (0.000)	0.0232 (0.000)	0.0402 (0.000)	0.0269 (0.000)
age_group	0.0006 (0.000)	0.0005 (0.220)	0.0075 (0.000)	0.0005 (0.299)	0.0013 (0.000)
gender	0.0001 (0.133)	0.0011 (0.000)	0.0005 (0.005)	0.0003 (0.078)	0.0001 (0.643)
education	0.0117 (0.000)	0.0050 (0.000)	0.0005 (0.003)	0.0012 (0.000)	0.0028 (0.000)
city_size	0.0007 (0.000)	0.0011 (0.001)	0.0020 (0.000)	0.0007 (0.065)	0.0019 (0.000)
econ_now	0.0004 (0.000)	0.0004 (0.173)	0.0011 (0.000)	0.0011 (0.000)	0.0010 (0.000)
econ_future	0.0027 (0.000)	0.0040 (0.000)	0.0009 (0.000)	0.0013 (0.000)	0.0023 (0.000)
econ_personal_future	0.0008 (0.000)	0.0027 (0.000)	0.0014 (0.000)	0.0008 (0.007)	0.0018 (0.000)
trust_people	0.0001 (0.004)	0.0002 (0.185)	0.0000 (0.791)	0.0012 (0.000)	0.0002 (0.207)
democracy_pref	0.0003 (0.000)	0.0033 (0.000)	0.0009 (0.000)	0.0019 (0.000)	0.0009 (0.001)
demo_satisfaction	0.0008 (0.000)	0.0005 (0.042)	0.0012 (0.000)	0.0010 (0.001)	0.0017 (0.000)
trust_parliament	0.0005 (0.000)	0.0008 (0.005)	0.0001 (0.722)	0.0009 (0.003)	0.0008 (0.003)
trust_gov	0.0010 (0.000)	0.0003 (0.223)	0.0004 (0.059)	0.0010 (0.001)	0.0012 (0.000)
trust_judiciary	0.0006 (0.000)	0.0007 (0.011)	0.0002 (0.363)	0.0007 (0.010)	0.0011 (0.000)
ideology_group	0.0009 (0.000)	0.0035 (0.000)	0.0011 (0.000)	0.0005 (0.057)	0.0021 (0.000)

Table 2: Mutual information (MI) between **fear_ai** and each feature, with p-values from Fisher's exact test (in parentheses). Age was discretized into quartiles prior to analysis. All MI values are rounded to 4 decimal digits, and p-values to 3.

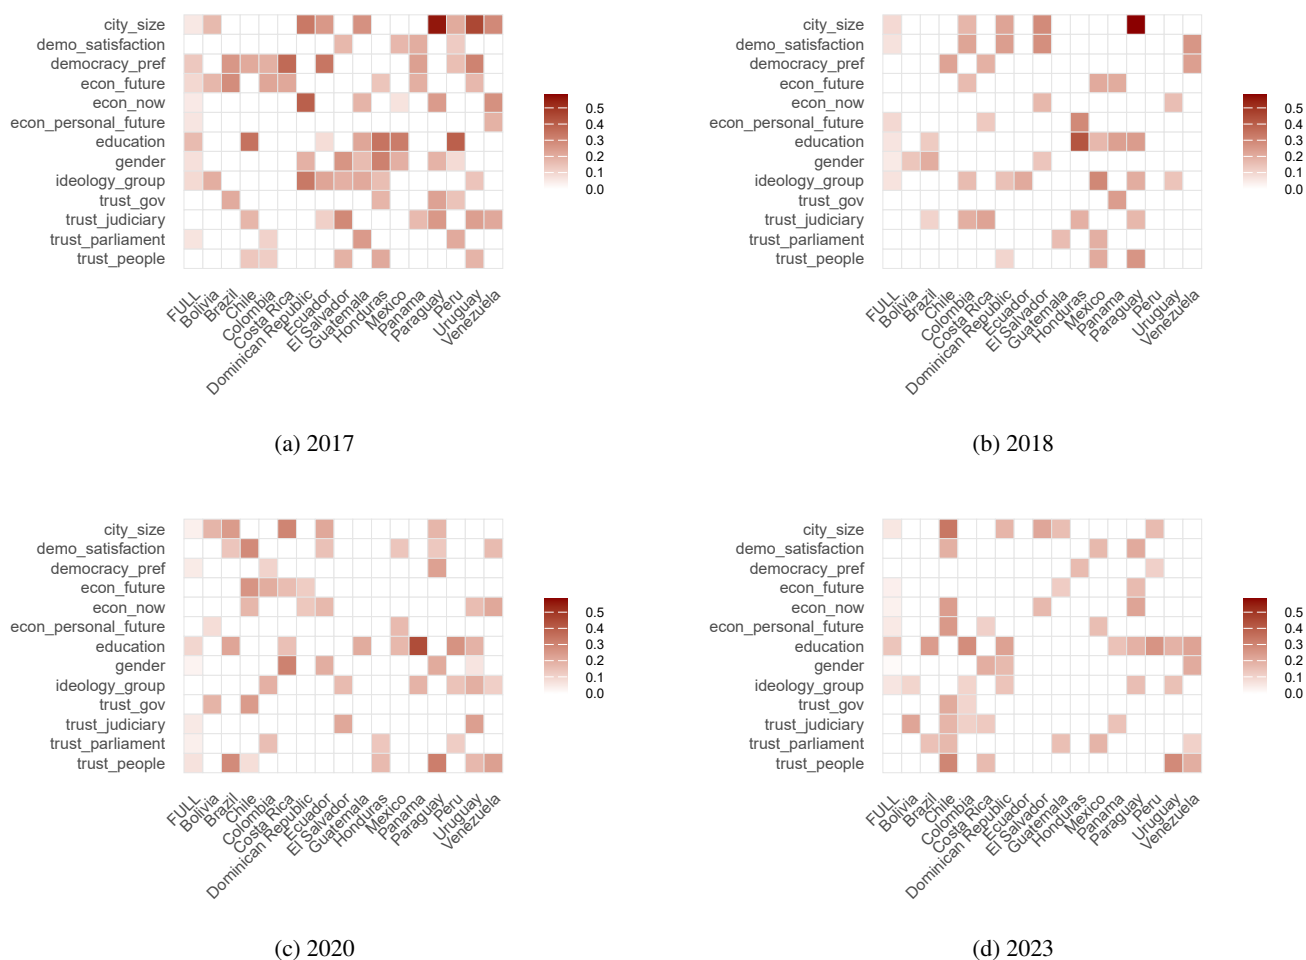


Figure 2: Variable inclusion and strength across countries for each year. White indicates the variable was excluded; darker shades represent higher average absolute coefficients.

trust_judiciary, suggesting that institutional confidence plays a particularly prominent role in Chilean attitudes toward automation. Colombia, on the other hand, presents a markedly different profile, with **ideology_group** dominating and **education** playing a secondary role, contrary to the broader regional trend.

Overall, the heatmaps show structural stability, such as the persistent influence of education and ideology, and cross-country differences in how social, economic, and political factors shape public fear of AI-driven job displacement.

Evolution of Key Predictors Over Time

Figure 3 displays the evolution of estimated coefficients for three key predictors, **education** (top), **political ideology** (middle), and **city size** (bottom), across countries and survey years. Each panel shows country-level estimates for each level of the predictor, with the pooled “FULL” model included for comparison. Because of the model structure and the use of sum-to-zero contrasts, negative coefficient estimates indicate greater fear of AI-driven job loss (i.e., a higher probability of agreeing or strongly agreeing that AI will lead to job displacement).

For **education**, individuals with lower attainment generally exhibit more negative coefficients, indicating a stronger association with fear of AI-driven job loss. Medium and high education levels show near-zero or slightly positive coefficients, consistent with a buffering effect of education.

For **political ideology**, coefficients are mostly close to zero, though the *Centre* category consistently shows slightly more positive values than Left or Right, suggesting that centrist respondents tend to express less concern about AI-driven job displacement. That said, the direction of coefficients for Left and Right varies by year: while left-leaning individuals show greater concern in earlier waves (e.g., 2017 and 2018), this pattern is less clear or even reversed by 2023.

For **city size**, respondents in large cities and capitals tend to show more positive coefficients, especially in 2023, indicating less fear of job loss in urban areas. In contrast, those in mid-sized cities and small towns typically show coefficients closer to zero or negative, pointing to a modest but persistent urban–rural divide.

Full-model estimates (available in the extended version) confirm these patterns, with education and ideology showing some of the largest coefficients across years and supporting their role as structural drivers of concern over AI-induced job loss.

The Role of Age in Shaping Fear of AI

Unlike other predictors, **age** is treated as a continuous variable in the ordinal regression models. It is included in approximately 32% of fitted models across countries and years, and when selected, the estimated coefficients are predominantly negative, suggesting that younger individuals tend to express greater concern about AI-driven job loss than older individuals. This negative association is most evident in 2017 and 2023, particularly in Bolivia, Costa Rica, El Salvador, and Peru (see the extended version).

The overall effect of age, however, is modest and varies across countries and years. In 2020, for instance, positive

coefficients emerge in Mexico, Paraguay, and Guatemala, indicating that in some contexts, older individuals express greater concern. These findings suggest that generational divides exist but are less consistent than those observed for education, political ideology, or city size, and likely reflect context-dependent patterns shaped by national or temporal factors.

Latent Class Analysis of Respondent Profiles

To identify distinct groups of respondents based on their demographic characteristics, economic perceptions, political attitudes, and trust in institutions, we perform a Latent Class Analysis (LCA) using all predictors except **country** and **fear_ai**. Although the BIC decreases monotonically with additional classes, likely due to the large sample size, models with more than four classes show diminishing separation and interpretability. We therefore retain a four-class solution as a compromise between BIC and entropy, balancing model fit and classification certainty (see the extended version for details).

Each class reflects a distinct combination of respondent characteristics, summarized below. Interpretations are based on the estimated item-response probabilities $\mathbb{P}(X_j = x_j \mid C = k)$, which represent the likelihood that individuals in latent class k choose a particular category of variable X_j . These probabilities inform the substantive meaning of each profile and are available in the extended version.

- **C1: Optimistic Institutionalists.** Respondents in this group tend to report high trust in government, parliament, and the judiciary, as well as relatively high satisfaction with democracy. They display strong economic optimism, particularly regarding future personal and national conditions, and are evenly distributed across age groups, genders, and education levels.
- **C2: Disillusioned Pessimists.** This group is characterized by overwhelmingly negative views of both the current and future economy, very low institutional trust, and high dissatisfaction with democracy. Respondents are somewhat older and more likely to live in small or mid-sized towns. They also report the highest generalized distrust in other people.
- **C3: Moderate Progressives.** These respondents tend to be younger, with above-average levels of education and moderate trust in institutions. They report moderate satisfaction with democracy and relatively good expectations for their personal economic future. Their ideological self-placement is more often centrist.
- **C4: Distrustful Urbanites.** This group includes respondents concentrated in large cities or capitals. They have low institutional trust, particularly in government and the judiciary, and a mixed economic outlook, more optimistic about their personal than national future. Their education levels are above average, and they endorse democracy at relatively high rates despite reporting dissatisfaction with its performance.

We begin by testing whether the distribution of **fear_ai** responses differs significantly across the four latent classes

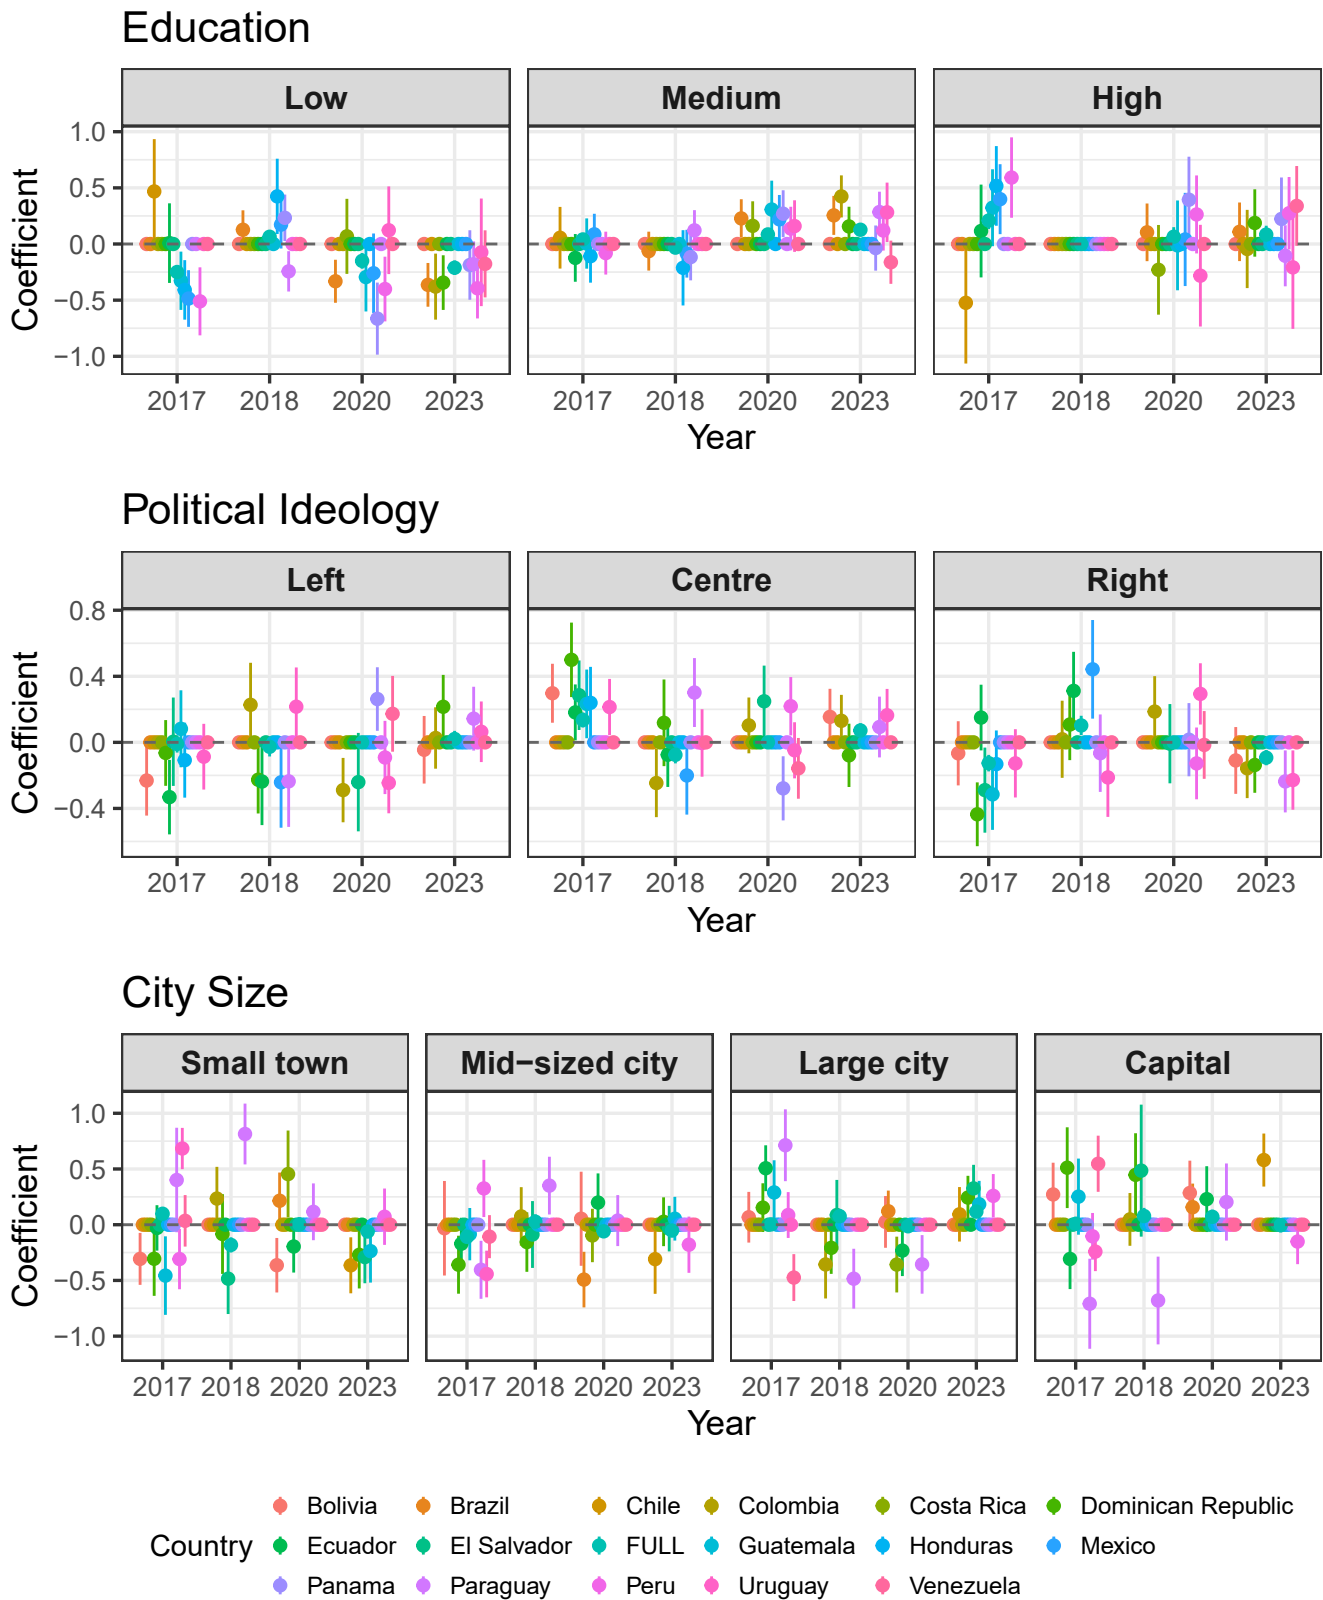


Figure 3: Evolution of estimated coefficients for **education** (top), **political ideology** (middle), and **city size** (bottom) across countries and years. Error bars indicate 95% confidence intervals.

Fear of AI	2017				2018				2020				2023			
	C1	C2	C3	C4	C1	C2	C3	C4	C1	C2	C3	C4	C1	C2	C3	C4
Strongly disagree	3%	5%	4%	3%	0%	0%	0%	0%	12%	14%	13%	12%	13%	15%	12%	14%
Disagree	21%	21%	21%	18%	18%	15%	17%	16%	45%	37%	42%	41%	43%	37%	40%	39%
Agree	48%	45%	50%	49%	69%	68%	70%	70%	31%	35%	34%	35%	30%	31%	34%	32%
Strongly agree	28%	29%	26%	30%	14%	17%	14%	14%	11%	13%	11%	12%	15%	17%	15%	16%

Table 3: Distribution of **fear_ai** responses (in %) across latent classes and survey years.

Class Pair	2017	2018	2020	2023
C1 vs C2	0.009	0.005	0.001	0.002
C1 vs C3	0.192	0.506	0.114	0.005
C1 vs C4	0.032	0.360	0.011	0.116
C2 vs C3	0.005	0.005	0.003	0.002
C2 vs C4	0.009	0.016	0.016	0.355
C3 vs C4	0.005	0.881	0.439	0.145

Table 4: Adjusted p -values (with Benjamini–Hochberg procedure) from pairwise chi-squared tests comparing **fear_ai** response distributions between latent classes, for each year.

within each survey year. Global chi-squared tests confirm strong evidence of association in all years (2017–2023), with p -values below 0.01 in every case. These results indicate that, although overall fear of AI is widespread, its expression varies systematically across respondent profiles.

To better understand these differences, Table 3 reports the percentage of responses in each fear category, by class and year. The differences in proportions appear modest, often just a few percentage points, but remain statistically significant in many comparisons, as shown in the pairwise tests reported in Table 4. For instance, in 2023, responses of “Strongly agree” range from 15% (C1) to 17% (C2), yet several class pairs show p -values below 0.01.

These subtle but robust differences align closely with the class profiles. **C2 (Disillusioned Pessimists)** consistently report the highest levels of concern about AI-induced job loss, reflecting their economic pessimism and institutional distrust. **C1 (Optimistic Institutionalists)** express the lowest concern, particularly in 2020 and 2023. **C3 (Moderate Progressives)** and **C4 (Distrustful Urbanites)** fall in between, but differ in political trust and urban exposure, with the latter group showing greater concern in more recent years.

To assess whether latent class membership captures attitudinal patterns that go beyond observed covariates, we examine how fear of AI-driven job loss varies with institutional trust across classes. Trust in government serves as a representative example, given its prominence in prior literature on AI perceptions. As shown in Figure 4, variation in fear levels is primarily explained by class membership rather than differences in trust alone. **C1 (Optimistic Institutionalists)** consistently reports the lowest concern, with a higher share of “Disagree” or “Strongly disagree” responses across all levels of trust. **C2 (Disillusioned Pessimists)**, by contrast, shows the highest fear regardless of reported trust. **C3 (Moderate Progressives)** and **C4 (Distrustful Urban-**

ites) occupy intermediate positions but differ in both baseline fear and sensitivity to trust. These patterns suggest that the latent classes capture stable attitudinal configurations not reducible to any single predictor.

Discussion

This study offers a comprehensive, multi-country analysis of how fear of AI-induced job loss has evolved in Latin America. Our findings reveal both widespread concern and notable heterogeneity in how these fears are structured across time, geography, and social groups.

Consistent with prior research in high-income countries, we find that education is the most consistent predictor of AI-related fear across time and countries (Bergdahl et al. 2023; Schiavo, Businaro, and Zancanaro 2024). Individuals with lower education levels are more likely to express concern about losing their job to AI and automation, while those with higher education levels report lower fear (Yarovenko et al. 2024). This pattern is evident in both the cross-country heatmaps and the full models, where education consistently exhibits some of the largest coefficients across years.

Political ideology plays a nuanced role in shaping attitudes toward AI. Across most years, respondents identifying with the political centre tend to express slightly less concern about AI-driven job loss compared to those on the ideological left or right. While earlier waves (e.g., 2017 and 2018) show somewhat greater fear among left-leaning individuals, this pattern becomes less consistent over time, with right-leaning respondents exhibiting greater concern in some countries by 2023. These findings partially mirror results from European studies such as Gerlich (2023), which highlight ideological divides in perceptions of automation, though in our data the gradient is modest and appears to shift over time rather than follow a fixed partisan pattern.

By contrast, the role of economic perceptions is more mixed. Although both national and personal economic outlooks are occasionally selected as predictors in the yearly models, their coefficients tend to be smaller and less consistent over time. This diverges from some findings in high-income contexts, where economic pessimism is a stronger driver of automation-related anxiety (Frey and Osborne 2017). Similarly, while institutional trust, particularly in government, parliament, and the judiciary, is included in some country-specific models, it plays a more secondary role in the pooled yearly estimates. However, the latent class results indicate that trust remains relevant for certain groups: **C2 (Disillusioned Pessimists)**, for instance, combine high concern about AI with persistent institutional distrust.




Figure 4: Distribution of **fear_ai** responses across levels of trust in government, stratified by latent class and survey year.

Latent Class Analysis further reveals that societal perceptions of AI cluster into distinct attitudinal profiles that extend beyond individual predictors. While all classes exhibit elevated concern in 2018, subsequent waves show more differentiated patterns. **C2 (Disillusioned Pessimists)** consistently express the highest fear, whereas **C1 (Optimistic Institutionalists)** show the lowest, particularly in 2020 and 2023. **C3 (Moderate Progressives)** and **C4 (Distrustful Urbanites)** fall in between, though they differ in their levels of institutional trust and urban exposure. These patterns suggest that fear of automation is shaped not only by observable traits but also by more stable attitudinal configurations that persist across survey years.

Overall, our findings show that fear of automation in Latin America cannot be reduced to a single explanation or confined to particular demographic categories. Rather, it reflects a dynamic interaction between education, political beliefs, economic expectations, and institutional trust. As AI technologies continue to advance, these fears may intensify or shift depending on how they intersect with broader trends in inequality, labor markets, and political identity.

Limitations

Several limitations should be considered when interpreting our findings. First, the data are observational and cross-sectional, which precludes causal inference. While we identify strong associations between predictors and fear of AI, we cannot determine the directionality of these relationships. Second, although we harmonized the outcome variable across survey waves, the wording of the question on AI-related job loss changed slightly over time. In particular, the 2018 item referenced the risk to both the respondent and their family, which likely contributed to the unusually high agreement observed in that year. Third, the statistical models rely on self-reported perceptions, which may be subject to reporting biases, particularly in politically sensitive environments. Finally, while our full models provide generalizable trends across the region, they may obscure mean-

ingful variation at the country level. For this reason, we included country-specific heatmaps and latent class profiles to highlight diversity in attitudes. Future work could incorporate panel data or event-based designs to capture shifts in response to economic or political developments.

Conclusions

This study contributes to the literature on public perceptions of AI by analyzing fear of job loss across four waves of the Latinobarómetro survey in 16 Latin American countries. Our findings highlight that concern about AI-driven automation is widespread, particularly among individuals with lower education levels and those identifying with the political left. Temporal and cross-country variation suggests that fear is shaped not only by individual characteristics but also by ideological context. Latent class analysis reveals four distinct respondent profiles, reflecting different combinations of demographic, economic, and political traits.

These findings have practical relevance for governments and organizations seeking to implement AI. Addressing public concern will require not only technical safeguards but also broader efforts to build institutional trust, reduce structural inequalities, and promote inclusive communication around the risks and benefits of automation. Policies aimed at upskilling the workforce, increasing transparency, and engaging citizens in AI governance may help mitigate fear and foster equitable technological adoption. These results underscore the need to address AI anxiety not as a technological issue, but as a social and political challenge shaped by education, trust, and inequality.

Acknowledgments

This paper was partially funded by PID2023-153222OB-I00 granted by MCIU / AEI / 10.13039/501100011033 / FEDER, UE.

References

- Acemoglu, D.; and Restrepo, P. 2020. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6): 2188–2244.
- Agresti, A. 2013. *Categorical data analysis*. John Wiley & Sons.
- Benjamini, Y.; and Hochberg, Y. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B*, 57(1): 289–300.
- Bergdahl, J.; Latikka, R.; Celuch, M.; Savolainen, I.; Mantere, E. S.; Savela, N.; and Oksanen, A. 2023. Self-determination and attitudes toward artificial intelligence: Cross-national and longitudinal perspectives. *Telematics and Informatics*, 82: 102013.
- Chiarini, A.; Grando, A.; Venturini, S.; and Borghonovo, E. 2023. Do automation and AI impact on job reduction? A study on perceived risk of losing job among white-collars in the Italian manufacturing companies. *Production Planning & Control*, 35(16): 2198–2211.
- Cremaschi, A.; Lee, D.-J.; and Leonelli, M. 2025. Understanding support for AI regulation: A Bayesian network perspective. *arXiv preprint arXiv:2507.05866*.
- Department for Science, Innovation and Technology. 2024. Public Attitudes to Data and AI: Tracker Survey (Wave 4).
- Dong, M.; Conway, J. R.; Bonnefon, J.-F.; Shariff, A.; and Rahwan, I. 2024. Fears about artificial intelligence across 20 countries and six domains of application. *American Psychologist*.
- Frey, C. B.; and Osborne, M. A. 2017. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114: 254–280.
- George, A. S. 2024. Artificial intelligence and the future of work: Job shifting not job loss. *Partners Universal Innovative Research Publication*, 2(2): 17–25.
- Gerlich, M. 2023. Perceptions and acceptance of artificial intelligence: A multi-dimensional study. *Social Sciences*, 12(9): 502.
- Gillespie, N.; Lockey, S.; and Curtis, C. 2021. Trust in Artificial Intelligence: A Five-Country Study. *University of Queensland / KPMG Report*.
- Gu, X.; Wang, G.; Obrenovic, B.; and Godinic, D. 2025. Fear of the new technology: investigating the factors that influence individual attitudes toward generative artificial intelligence. *Current Psychology*.
- Kaya, F.; Aydin, F.; Schepman, A.; Rodway, P.; Yetişensoy, O.; and Kaya, M. D. 2024. The roles of personality traits, AI anxiety, and demographic factors in attitudes toward artificial intelligence. *International Journal of Human-Computer Interaction*, 40(2): 497–514.
- Latinobarómetro Corporation. 2017–2023. Latinobarómetro Survey Dataset. <https://www.latinobarometro.org>. Waves used: 2017, 2018, 2020, 2023. Accessed 2024.
- Lazarsfeld, P. F. 1950. The Logical and Mathematical Foundations of Latent Structure Analysis. In Stouffer, S. A., ed., *Measurement and Prediction*, 362–412. New York: John Wiley & Sons.
- Li, J.; and Huang, J.-S. 2020. Dimensions of artificial intelligence anxiety based on the integrated fear acquisition theory. *Technology in Society*, 63: 101410.
- Linzer, D. A.; and Lewis, J. B. 2011. poLCA: An R package for polytomous variable latent class analysis. *Journal of Statistical Software*, 42: 1–29.
- Lund, B. D.; Mannuru, N. R.; and Agbaji, D. 2024. AI anxiety and fear: A look at perspectives of information science students and professionals towards artificial intelligence. *Journal of Information Science*, 1–9.
- McCullagh, P. 1980. Regression models for ordinal data. *Journal of the Royal Statistical Society Series B*, 42(2): 109–127.
- Meyer, P. E. 2022. *infotheo: Information-Theoretic Measures*. R package version 1.2.0.1.
- Montag, C.; Kraus, J.; Baumann, M.; and Rozgonjuk, D. 2023. The propensity to trust in (automated) technology mediates the links between technology self-efficacy and fear and acceptance of artificial intelligence. *Computers in Human Behavior Reports*, 11: 100315.
- Morikawa, M. 2017. Who are afraid of losing their jobs to artificial intelligence and robots? Evidence from a survey. *GLO Discussion Paper Series*, 71.
- Nnamdi, N.; Ogunlade, B. Z.; and Abegunde, B. 2023. An evaluation of the impact of artificial intelligence on socio-economic human rights: A discourse on automation and job loss. *Scholars International Journal of Law, Crime and Justice*, 6(10): 508–521.
- Posit team. 2024. *RStudio: Integrated Development Environment for R*. Posit Software, PBC, Boston, MA.
- Schiavo, G.; Businaro, S.; and Zancanaro, M. 2024. Comprehension, apprehension, and acceptance: Understanding the influence of literacy and anxiety on acceptance of artificial intelligence. *Technology in Society*, 77: 102537.
- Soueidan, M. H.; and Shoghari, R. 2024. The impact of artificial intelligence on job loss: Risks for governments. *Technium Social Sciences Journal*, 57: 206–223.
- Uçar, M.; Çapuk, H.; and Yiğit, M. F. 2025. The relationship between artificial intelligence anxiety and unemployment anxiety among university students. *Work: A Journal of Prevention, Assessment and Rehabilitation*, 80(2): 701–710.
- Venables, W. N.; and Ripley, B. D. 2002. *Modern applied statistics with S*. New York: Springer, fourth edition.
- Wickham, H. 2016. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York.
- Wilde, N. 2025. Fear of artificial intelligence or fear of looking in the mirror? Revisiting the Western machine-takeover imaginary. *AI & Society*, 1–11.
- Yarovenko, H.; Norek, T.; Kuzior, A.; and Lopatka, A. 2024. The future of artificial intelligence: Fear, hope or indifference? *Human Technology*, 20(3): 611–639.
- Zhan, E. S.; Molina, M. D.; Rheu, M.; and Peng, W. 2024. What is there to fear? Understanding multi-dimensional fear

of AI from a technological affordance perspective. *International Journal of Human–Computer Interaction*, 40(22): 7127–7144.

Zhang, B.; and Dafoe, A. 2019. Artificial intelligence: American attitudes and trends. *Center for the Governance of AI, Future of Humanity Institute*.