

WORKING PAPER IDB-WP-1832

Regulating AI Displacement

Public Demand, Institutional Credibility, and the Politics of Employment Protection in Latin America

Maria Victoria Rosino
Carlos Scartascini

Inter-American Development Bank
Department of Research and Chief Economist

October 2026



Regulating AI Displacement

Public Demand, Institutional Credibility, and the Politics of Employment Protection in Latin America

Maria Victoria Rosino[†]

Carlos Scartascini[‡]

[†] Universidad de San Andrés

[‡] Inter-American Development Bank

Inter-American Development Bank
Department of Research and Chief Economist

October 2026



JEL Codes: D72, D83, J38, O33

Keywords: AI governance, regulation, employment protection, trust, survey experiment

To explore more IDB publications, visit

<https://publications.iadb.org/en>

Copyright © 2026 Inter-American Development Bank (IDB). This work is subject to a Creative Commons Attribution 4.0 International Public License CC BY 4.0 (<https://creativecommons.org/licenses/by/4.0/legalcode>). The terms and conditions indicated in the URL link must be met and the respective recognition must be granted to the IDB.

Any and all disputes arising under this license that cannot be settled amicably shall be resolved in accordance with the following procedure. Pursuant to a notice of mediation communicated by reasonable means by either you or the licensor to the other, the dispute shall be submitted to non-binding mediation conducted in accordance with the World Intellectual Property Organization (WIPO) Mediation Rules. Any dispute that cannot be settled amicably shall be submitted to arbitration pursuant to the United Nations Commission on International Trade Law (UNCITRAL) rules. The use of the IDB name for any purpose other than the respective recognition and the use of the IDB logo are not authorized by this license and require an additional license agreement.

Note that the URL link includes terms and conditions that are an integral part of this license.

The opinions expressed in the work are those of its authors and do not necessarily reflect the views of the IDB, its Board of Executive Directors, or the countries they represent.



Abstract*

How citizens want artificial intelligence (AI) regulated will shape which AI governance instruments are politically feasible. Yet we know little about public preferences over AI policy in middle-income democracies, where state capacity is contested. This paper provides the first multi-country experimental evidence on preferences over AI labor-displacement policy instruments in middle-income democracies. Using a randomized framing experiment in a survey of Argentina, Brazil, Colombia, and Mexico, we examine how citizens map beliefs about AI-driven job disruption into demand for three governance responses: worker-retention regulation, tax-financed redistribution to displaced workers, or non-intervention. Nearly three-quarters of respondents prefer employment-preserving regulation, about one in twenty favors tax-financed redistribution, and one-fifth prefer non-intervention. A three-arm framing taxonomy spanning alarmist, delayed-threat, and task-reallocation narratives produces no average shift in this distribution, and the heterogeneous effects we detect at the margin run from non-intervention toward regulation, not toward redistribution. The pattern is consistent with preferences anchored in priors about the state’s redistributive capacity rather than in correctable beliefs about AI, with direct implications for AI governance design in low-credibility institutional environments. A within-survey wording experiment on employment protection regulation elicits a substantial shift in preferences among the same respondents, indicating that the null on the AI framings is not an artifact of respondent inattention.

* The authors gratefully acknowledge LAPOP at Vanderbilt University’s collaboration for data collection. Data collection was IRB-approved by the Behavioral Sciences Committee of Vanderbilt University’s IRB. Pre-trial registration available at: AEARCTR-0016600. The Inter-American Development Bank (IDB) provided funding. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported herein. IDB management had no involvement in the study design, analysis, or interpretation of the data; in the writing of the report; or in the decision to submit the article for publication. The opinions expressed here are those of the authors and do not necessarily reflect the views of the IDB, its Board of Directors, or the countries they represent. This work was created with assistance from Claude Sonnet 5 to check grammar, style, spelling, and punctuation. A professional editor made a final copyedit.

1 Introduction

Artificial intelligence (AI) is generating an active regulatory debate. Some governments are designing labor-market protections, retraining programs, competition rules, and tax redistribution schemes in parallel, often under uncertainty about how the technology will reshape work.¹ In most countries, which of these instruments proves politically feasible will depend not only on their welfare properties but also on what citizens actually demand and on whether that demand can be reshaped by information about the technology itself. This paper asks what citizens want their governments to do about AI-driven labor displacement and how those preferences shift when challenged with new information about AI’s impact on employment.

We bring evidence to bear on this question using a randomized framing experiment embedded in an online survey conducted in 2025 in Argentina, Brazil, Colombia, and Mexico ($N = 3,996$). Respondents are randomly assigned to a control group and to informational vignettes (actual newspaper headlines from Latin America) that frame AI-induced job loss as large and immediate (*alarmist*), as delayed and gradual (*delayed threat*), or as removing low-value work (*task reallocation*). We then measure support for three governance responses: regulations requiring firms to retain workers, taxes on the employed to fund transfers to the displaced, or non-intervention.² The design isolates two distinct questions: what citizens demand at baseline (control group) and whether that demand responds to informational variation in how AI risk is framed (comparisons between treatment and control).

Three findings emerge. First, the baseline distribution of preferences is highly skewed. Nearly three-quarters of respondents favor worker-retention regulation, roughly one in twenty favors tax-financed redistribution, and about one-fifth prefer no government intervention. The pattern holds across all four countries, varies modestly in level, and follows a clear ideological gradient and perceived job exposure to AI. Second, the framing treatments produce no statistically distinguishable average shift in this distribution. Whether respondents are told that AI could eliminate up to 60 percent of current jobs, that major effects are a decade away and small, or that AI will mainly remove low-value tasks, their aggregate policy preferences barely move. Third, heterogeneous effects exist at the margin: lower-education respondents, the unemployed, and those on the left of the political ideology spectrum are somewhat more responsive to framing, but even for these groups, movement runs from non-intervention toward retention regulation rather than

¹For example, the Governor of California signed an executive order “to prepare workers and businesses for potential AI disruption” (Executive Department, State of California, 2026). Several major bills have been introduced in the U.S. Congress to address the impact of artificial intelligence on jobs, workforce data, and the economy. For example, the “AI Tax and Work Protection Act” proposes an excise tax tied to foundation-model use to finance a Work Protection Administration and employment programs responding to AI-related displacement (U.S. House of Representatives, 2026).

²These alternatives summarize well part of the current debate on AI regulation, which fluctuates between “how to reduce job losses” (e.g., Bill Gates’ view (Svirnovskiy, 2026)) and “checks issued by the Federal government [are]... the best way to deal with unemployment caused by AI” (e.g., Elon Musk’s view (Siddiqui, 2026)).

toward redistribution. Perceived AI-job-vulnerability does not matter significantly.

To test whether the null result is an artifact of uniform respondent inattention, we also use a within-survey wording experiment on employment protection regulation, administered on the same instrument to the same respondents. This within-subject comparison elicits a substantial and highly significant shift in expressed preferences. Therefore, preferences do move when the manipulation touches a conceptually adjacent labor-policy question.

The experimental null is itself informative for the politics of AI regulation. Because the treatments span a wide range of credible narratives, the absence of average movement is inconsistent with the interpretation that the gap between citizen demand and the broader portfolio of welfare-oriented responses emphasized in the economics literature—including redistribution, tax reform, competition policy, and policies steering technological development (Acemoglu & Restrepo, 2018; Acemoglu et al., 2026; Athey & Scott Morton, 2026; Guerreiro et al., 2022; Korinek & Lockwood, 2026; Korinek & Stiglitz, 2025)—reflects specific, correctable misbeliefs about AI. The preference for protection appears to reflect deeper priors that short informational interventions cannot move and is compatible with work on the politics of regulation under weak state capacity: when citizens doubt the state’s ability to follow through on promised compensation, they tend to demand protective regulation that delivers visible insurance rather than fiscal instruments requiring credible delivery (Keefer & Scartascini, 2022; Levy & Spiller, 1994, 1996). Latin America’s persistently low levels of institutional trust (Keefer & Scartascini, 2022; Keefer et al., 2022) are consistent with the near-absence of demand for tax-financed redistribution we observe. The implications for AI governance are direct: if citizen demand for employment protection is not contingent on beliefs about AI’s specific impacts, communication-based legitimation strategies (informational campaigns, expert reports, deliberative consultations centered on the technology) may be unlikely to shift demand toward instruments preferred on welfare grounds.

The paper contributes to three literatures. To scholarship on AI governance, where survey evidence has established broad public support for regulating AI but has been gathered almost entirely in high-capacity democracies (Laux et al., 2024; Lundgren & Tallberg, 2026; Zhang & Dafoe, 2020), we provide systematic evidence on public preferences over AI regulatory instruments in middle-income democracies, documenting a sharp gap between citizen demand and the policy portfolio that the welfare economics literature identifies as efficient. To work on the politics of regulation in developing countries, where the state’s limited capacity to commit shapes which instruments citizens find credible (Braithwaite, 2006; Dussauge-Laguna, 2024; Levy & Spiller, 1994, 1996; Minogue & Cariño, 2006), we identify a setting in which informational legitimation appears ineffective, speaking to debates about when informational and deliberative strategies can move regulatory preferences (Capelos et al., 2016; Maman et al., 2025). The pattern we document is one of social protection demanded through regulatory rules rather than fiscal transfers, the

arrangement scholars of regulatory governance call a *regulatory welfare state* (Benish & Levi-Faur, 2020; Horton & Lynch-Wood, 2024): the instrument citizens prefer is the one whose enforcement burden falls on firms rather than on the fiscal authority. We further demonstrate, through a within-survey wording experiment, that this ineffectiveness is specific to the AI framing rather than general to the sample or instrument. To the literature on automation and policy attitudes (Gallego et al., 2022; Golin & Rauh, 2025; Jeffrey, 2021; Jeffrey & Matakos, 2024; Lee, 2024; Magistro et al., 2024; van Hoorn, 2022), we add evidence from a region where prior surveys have shown that information can shift preferences on other contentious policy domains (Cafferata et al., 2023; Domínguez & Scartascini, 2024; Gingerich & Scartascini, 2022), which makes the AI-specific null particularly informative. One specific comparison anchors this: Duval et al. (2024) find that information treatments shift employment-protection-legislation preferences by up to 13 percentage points in a cross-country experiment of developed countries, while we find effects indistinguishable from zero. Together with the Latin American comparison, the null appears specific to the AI context and region, rather than to the survey instrument or to regulatory attitudes more generally.

2 Data and Sample Composition

The survey targets 1,000 adult residents (18 years and older) in each of four Latin American countries: Argentina, Brazil, Colombia, and Mexico. It was administered online in 2025 through the Latin American Public Opinion Project (LAPOP). The questionnaire includes an embedded information-framing experiment on AI and employment. Respondents are randomly assigned (with equal probability) to one of four conditions: a control group that receives only a baseline definition of AI, or one of three treatments that frame AI’s labor-market impact in distinct ways: an *Alarmist* frame emphasizing large potential job displacement, a *Delayed Threat* frame emphasizing gradual and distant impacts, and a *Job Reallocation* frame emphasizing the elimination of low-value work through automation. After exposure, respondents report their preferred policy response to a scenario in which AI reduces labor demand. Table 1 summarizes the content of each treatment; full wording appears in Online Appendix A. Together, these four countries comprise a substantial share of Latin America’s adult population and share the region’s low levels of institutional trust, while differing in labor-market structure and recent regulatory trajectories (Keefer & Scartascini, 2022).

The treatments are based on real newspaper headlines from El País and Infobae, which have large regional readership (PRISA, 2025; WashPostPR, 2020). The articles draw on different sources: PwC, McKinsey, and the World Economic Forum in one case; Goldman Sachs in another; and The New York Times in the third. Using actual headlines exposes respondents to information in a realistic and credible format but at the cost of less control over message content. The treatments, therefore, capture bundled exposure to real media rather than clean semantic framing effects.

We flag one limitation upfront: no post-treatment manipulation check was administered to verify whether respondents processed the framing content. The null average treatment effects we report below are therefore consistent with two interpretations: that preferences are robust to information about AI risk or that respondents engaged only lightly with the headlines. We return to this design limitation and its implications for the interpretation of the null result later, when we draw on a within-survey wording experiment on employment protection regulation to bound the weak first-stage interpretation.

Table 1: Treatment arms: Framing of AI’s labor-market impact

<i>Panel A: Treatment Assignment</i>	
Control ($p \sim 1/4$)	Respondents receive the baseline definition of AI and proceed directly to the policy question.
Alarmist Framing ($p \sim 1/4$)	Respondents read an actual newspaper headline reporting estimates that AI-driven automation could eliminate <i>up to 60%</i> of current jobs.
Delayed Threat ($p \sim 1/4$)	Respondents read an actual newspaper headline emphasizing gradual impacts: only about 5% of tasks are currently automatable and major labor-market effects are not expected for at least <i>a decade</i> .
Job Reallocation ($p \sim 1/4$)	Respondents read an actual newspaper headline stating that AI could <i>eradicate “meaningless” jobs</i> by automating tasks that many consider useless.
<i>Panel B: Experimental Outcome</i>	
Policy preference if AI reduces demand for labor ^a	(1) Regulations requiring companies to hire workers and preserve employment; (2) Increasing taxes on those who work to subsidize the unemployed; (3) Not introducing regulations or increasing taxes. ^b

Notes: The full wording of the headlines and the outcome question is reproduced in Online Appendix A.^a A reduction in the demand for labor can occur through channels beyond a reduction in the number of jobs, including fewer hours, lower wages, and reduced fringe benefits; the question asks about the quantity of employment to provide respondents a salient and concrete reference.^b The three options map onto the generic government responses to displacement risk: mandating employment (protection), financing compensation through taxation (redistribution), and allowing market adjustment (non-intervention). This menu does not include other instruments such as capital taxation, competition enforcement, or innovation subsidies, which are prominent in the welfare literature but less salient as policy options in public discourse in the region.

The survey yielded $N = 3,996$ completed responses. Of these, 3,957 have non-missing values on the policy-preference outcome and form our analytic sample; the 39 excluded respondents are distributed roughly evenly across treatment arms and countries. Assignment is approximately even across arms, and country-level sample sizes are also well-balanced. The sample skews toward higher education: approximately 60% of respondents have more than a high school degree, which is

above the share in the adult population of these countries. Results should therefore be interpreted as reflecting the preferences of a more educated urban segment rather than the full electorate. The employment variable used in the analysis captures current employment status (employed vs. not employed). LAPOP identifies formal private-sector employees under contract and public-sector employees but does not separately identify formal from informal own-account workers and small business owners, which limits a full formality-based heterogeneity analysis.

Table 2 reports the distribution of assignment by country; the cell counts are nearly identical across arms. Throughout the analysis, we measure stated policy preferences in a survey context; whether these translate into actual voting behavior, lobbying, or policy-change pressure is a separate question that we cannot address with the available data.

Table 2: Distribution of framing treatment assignment by country

	Argentina	Brazil	Colombia	Mexico	Total
<i>Control</i>	251	251	248	251	1001
<i>Alarmist</i>	246	250	251	252	999
<i>Delayed Threat</i>	251	249	251	250	1001
<i>Job Reallocation</i>	251	249	250	245	995
Total	999	999	1000	998	3996

Table B.1 in Online Appendix B summarizes baseline respondent characteristics by experimental arm, reporting group means (and standard deviations) for a broad set of pre-treatment covariates, including demographics (gender, age), socioeconomic indicators (education, wealth, employment), pre-treatment AI beliefs and perceived exposure to job displacement, risk-related measures, proxies for civic attitudes, and basic economic knowledge.³ Overall, the randomization appears to have achieved balance: most covariates do not differ statistically across groups. A small number of pairwise differences relative to the Control group reach conventional significance levels (e.g., modest differences in the female share, employment status, and some economic-knowledge items), but these are limited in number and small in magnitude, consistent with chance variation under multiple testing. Taken together, these patterns support interpreting subsequent differences in outcomes across framing arms as causal intent-to-treat effects.

3 Baseline Policy Preferences and Beliefs about AI Displacement

A large majority of respondents favor worker-retention regulation (72.3% in the control group) when asked what policy they prefer if AI reduces labor demand. Only a small minority endorses higher

³Online Appendix A reproduces the survey questions and documents the coding of each variable in Table A.1.

taxes on employed workers to subsidize the unemployed (5.8%), while 21.9% prefer no government intervention. Table 3 disaggregates these shares by respondent characteristic, restricted to the control group to avoid treatment contamination.

Table 3: Baseline distribution of policy support preferences by respondent characteristics

	<i>Policy support preference (%)</i>		
	Worker-retention regulation	Tax-financed redistribution	No government intervention
Overall	72.3	5.8	21.9
Country			
Argentina	65.7	3.6	30.7
Brazil	71.7	8.0	20.3
Colombia	78.6	3.6	17.7
Mexico	73.3	8.0	18.7
Education			
Less than high school	65.5	5.5	29.1
Completed high school	73.5	7.1	19.4
More than high school	73.0	5.2	21.8
Political ideology			
Left	77.4	6.6	16.0
Center-left	78.2	8.0	13.8
Center	76.0	4.5	19.6
Center-right	66.0	3.8	30.2
Right	59.7	8.4	31.8
Perceived personal displacement risk from AI			
AI could not replace my job	70.2	4.7	25.1
AI could replace my job	75.0	7.1	17.9
Perceived incidence of AI-related unemployment across social groups			
Poor	75.3	6.0	18.7
Middle class	70.2	5.2	24.6
Upper class	64.1	10.3	25.6

Notes: Shares of respondents in the control group selecting each policy option: worker-retention regulation, tax-financed redistribution, or no government intervention. Percentages sum to 100 within each row. Perceived personal displacement risk from AI is measured by respondents' self-assessment of whether AI could replace their job. Perceived incidence of AI-related unemployment reflects the social group (poor, middle class, or upper class) that respondents identify as most affected by AI-driven job displacement.

The comparatively low support for tax-financed redistribution could in principle reflect the specific cost architecture of the choice set, namely, that only employed respondents would bear the proposed tax. We test this directly on the control arm (Online Appendix Table B.2). Baseline support for redistribution is essentially flat across labor-market positions: 5.96% among working respondents, 4.35% among the unemployed, and 5.88% among those out of the labor force. A linear regression

with country fixed effects yields coefficients of -0.016 ($p = 0.50$) for unemployed and -0.005 ($p = 0.79$) for out-of-labor-force respondents relative to workers as the omitted category. Within the subset of workers, a suggestive pattern by formality appears: formal workers report 4.81% support versus 7.69% for informal and 7.18% for self-employed (the joint test of equal support fails to reject ($F(2, 629) = 1.00$, $p = 0.37$)). The low aggregate support therefore reflects a cross-labor-market aversion rather than a mechanical response to who would bear the tax proposed.

The distribution of responses reveals a sharp gap between the regulatory instruments citizens demand and the broader policy portfolios advanced in expert AI-governance debates, which tend to emphasize redistribution, retraining support, and competition policy over employment mandates (Figure B.1 summarizes the responses shown in Table 3). This pattern of demand is compatible with evidence from low-trust settings like Latin America, where citizens instead prefer regulatory instruments because they deliver protection through visible, non-discretionary mechanisms that do not depend on the government’s future implementation capacity. Employment mandates succeed because firms bear the enforcement cost; redistribution requires sustained government credibility across time. Recent experimental evidence from pension-reform debates in the region confirms this pattern (Bonvecchi et al., 2026).

The low preference for tax-financed redistribution is consistent with the well-documented low levels of institutional trust across the region (Keefer & Scartascini, 2022; Keefer et al., 2022). One interpretation is that when citizens doubt the state’s capacity to tax and deliver transfers effectively, visible employment protection may serve as a more reliable form of insurance, though we cannot test this mechanism directly with the cross-sectional data available here.

Preference ordering is similar across all four countries for the control group. Support for worker-retention regulation ranges from 65.7% in Argentina to 78.6% in Colombia, with Brazil (71.7%) and Mexico (73.3%) in between. Differences across countries are driven mainly by the size of the “no intervention” group, which is largest in Argentina and smallest in Colombia, while the redistributive option remains comparatively uncommon in every context. Policy distribution by education suggests modest compositional differences: support for worker-retention regulation is high across all education groups, but somewhat lower among respondents with less than high school education, mirrored by a larger “no intervention” share in that group.

Political ideology is also strongly associated with the preferred policy response. Support for worker-retention regulation remains the modal choice, but it is systematically lower among respondents who place themselves on the right of the political spectrum, where the “no intervention” option becomes comparatively more prevalent. By contrast, respondents on the left are substantially more likely to favor regulatory protection and least likely to endorse a laissez-faire response.

Understanding public support for labor-market regulation in response to AI requires going beyond

average attitudes and considering perceived heterogeneity in exposure to technological change. Individual beliefs about replaceability can proxy expected adjustment costs and may therefore be closely associated with the policy response respondents deem appropriate. To capture this dimension, the survey asks respondents (before the randomized section) whether they believe that artificial intelligence could replace their job. Figure B.3 in Online Appendix B shows that 44.9% of respondents report that AI could replace their job, compared with 55.1% who report it could not, with a clear education gradient.

Respondents who view themselves as replaceable are more likely to favor worker-retention regulation (75.0% versus 70.2%) and less likely to support no government intervention (17.9% versus 25.1%), alongside a modest increase in support for the redistributive option. While descriptive, this pattern is consistent with a political-economy linkage between perceived labor-market risk and the preferred scope of government involvement in the adjustment process.

Beliefs about the distributional incidence of AI-driven displacement also shape policy preferences. Respondents overwhelmingly locate AI-related unemployment risk among lower- and middle-income groups (Figure B.5 in Online Appendix B). Those who expect displacement to concentrate among the better-off show weaker support for worker-retention regulation, stronger support for no government intervention, and a higher propensity to favor the redistributive option (10.1%, versus 5.5% in the full sample; Figure B.6 in Online Appendix B). This pattern relates to Jeffrey and Matakos (2024), who show that fairness beliefs and prior exposure to labor-market shocks condition redistributive responses to automation risk.

4 Empirical Strategy

We estimate intent-to-treat (ITT) effects using a linear probability model. Let Y_i indicate the policy option chosen by respondent i . The baseline specification is:

$$Y_i = \alpha + \beta_1 Alarmist_i + \beta_2 Delayed_i + \beta_3 Reallocation_i + \mathbf{X}_i' \boldsymbol{\delta} + \lambda_{c(i)} + \varepsilon_i, \quad (1)$$

where $Alarmist_i$, $Delayed_i$, and $Reallocation_i$ are binary indicators for treatment assignment. The control group serves as the omitted category. The coefficients β_1 through β_3 capture the average treatment effects of each framing relative to the control group. The vector $\mathbf{X}_i$ includes pre-treatment covariates: age, gender, education indicators, wealth index, employment status, political ideology, personal displacement risk, AI vulnerability group, perceived collective action, state responsiveness, perceived police-response efficiency, patience, risk tolerance, measures of basic economic knowledge, and a news-following indicator.⁴ λ_c denotes country fixed effects and ε_i is an individual error term.

⁴Detailed coding of each variable from the original survey items is reported in Online Appendix Table A.1.

Because the outcome consists of three unordered categories, we also estimate a multinomial logit model as a robustness check, using *no government intervention* as the base outcome. The specification and estimates are reported in Online Appendix C.

To examine whether average effects mask offsetting responses across subgroups, we interact treatment assignment with six pre-specified respondent characteristics: education, employment status, perceived AI vulnerability group, age, political ideology, and gender. For each policy outcome indicator $y_{i,k} \in \{0, 1\}$, the interacted specification is:

$$y_{i,k} = \alpha_k + \sum_{t \in \mathcal{T}} \beta_{t,k} T_{i,t} + \sum_{g \in \mathcal{G} \setminus \{g_0\}} \gamma_{g,k} \mathbb{1}\{G_i = g\} + \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G} \setminus \{g_0\}} \delta_{t,g,k} T_{i,t} \times \mathbb{1}\{G_i = g\} + \mathbf{X}_i' \theta_k + \lambda_{c(i),k} + \varepsilon_{i,k}, \quad (2)$$

where $T_{i,t}$ are the treatment indicators, G_i denotes the heterogeneity dimension of interest (with baseline category g_0 omitted), $\mathbf{X}_i$ is the full covariate set, and $\lambda_{c(i),k}$ are country fixed effects. The interaction terms $\delta_{t,g,k}$ capture how the treatment effect in group g differs from the effect in the baseline group g_0 .

5 Results

5.1 Baseline treatment effects on policy preferences

Table 4 reports ITT estimates from the baseline linear probability model (Equation 1).⁵ Across all three outcomes, framing effects are small and statistically indistinguishable from zero. A joint Wald test on the treatment coefficients across the three outcomes fails to reject the null of no framing effect ($\chi^2(6) = 2.13$, $p = 0.907$).⁶ This result carries a specific interpretation: because the treatments span a wide range of narratives, from forecasts of large immediate job loss to accounts of gradual task reallocation, the absence of movement suggests that preferences for employment protection are not anchored in particular beliefs about the scale or timing of AI disruption that new information could correct. The null contrasts with Duval et al. (2024), who find that information treatments shift EPL preferences by up to 13 percentage points in a related cross-country survey experiment, as well as with Latin American evidence showing that informational treatments do shift preferences on other contentious topics (Cafferata et al., 2023; Domínguez & Scartascini, 2024; Gingerich & Scartascini, 2022).

We quantify this contrast formally through minimum detectable effect (MDE) calculations and equivalence tests (Table 5). At 80% power ($\alpha = 0.05$), the design can detect effects as small as 5.4 percentage points on worker-retention regulation, 2.9 percentage points on tax-financed

⁵Full coefficient estimates, with and without covariate adjustment, are reported in Online Appendix Table B.3.

⁶The joint test is computed across the three linear probability models with robust standard errors. The six degrees of freedom reflect the linear dependency among coefficients across mutually exclusive outcomes.

redistribution, and 5.0 percentage points on no government intervention. Equivalence tests using the two one-sided test (TOST) procedure at the 5% level reject the null of an effect of ± 13 percentage points (the magnitude reported by Duval et al. (2024)) for all nine treatment-outcome combinations. Against a more stringent threshold of ± 5 percentage points, eight of nine tests still reject; the only exception is the Delayed Threat treatment on worker-retention regulation, where the upper bound of the 90% confidence interval equals 0.052.

Table 4: Treatment effects on policy preferences (OLS estimates)

	<i>Policy preference</i>		
	Worker-retention regulation (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	0.010 (0.019)	0.002 (0.010)	-0.012 (0.018)
Delayed Threat	0.020 (0.019)	-0.009 (0.010)	-0.011 (0.018)
Job Reallocation	0.006 (0.019)	-0.003 (0.010)	-0.002 (0.018)
Personal displacement risk	0.029** (0.014)	0.020*** (0.007)	-0.049*** (0.013)
AI vulnerable group: Middle	-0.012 (0.015)	0.003 (0.008)	0.009 (0.014)
AI vulnerable group: Upper	-0.091** (0.037)	0.053** (0.024)	0.039 (0.033)
Constant	0.644*** (0.089)	0.183*** (0.051)	0.173** (0.081)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	3,957	3,957	3,957
R-squared	0.059	0.032	0.068
Joint significance (p -value)	0.774	0.723	0.865
Within-source equality (p -value)	0.810	0.597	0.569
Between-source equality (p -value)	0.475	0.340	0.807

Notes: Each column reports the regression coefficients and robust standard errors in parentheses from OLS estimations of linear probability models. The dependent variable is an indicator equal to 1 if the respondent selects the corresponding policy option and 0 otherwise: worker-retention regulation (Column 1), redistribution through higher taxes on workers to subsidize the unemployed (Column 2), and no government intervention (Column 3). Coefficients on treatment indicators (Alarmist, Delayed, Job Reallocation) capture ITT effects relative to the Control group (omitted category). Joint significance is the F-test of the null that the three treatment coefficients (Alarmist, Delayed Threat, Job Reallocation) jointly equal zero within each outcome. Within-source equality tests the null that the Alarmist and Job Reallocation coefficients (both from Infobae) are equal within each outcome. Between-source equality tests the null that the average of the two Infobae coefficients equals the El País (Delayed Threat) coefficient within each outcome. All columns include the full set of pre-treatment covariates and country fixed effects. Controls include gender, age and age squared, education indicators, wealth index, employment status, political ideology, collective-action measures, perceived police-response efficiency, patience, risk tolerance, basic economic-knowledge measures, a news-following indicator, personal displacement risk, and AI vulnerability group. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The three treatments are drawn from two different newspapers (Alarmist and Job Reallocation from Infobae, Delayed Threat from El País) so source variation and framing variation are partly confounded across arms. Two tests bound this confounding. The two Infobae-sourced treatments yield coefficients that are statistically indistinguishable within each outcome (within-source p -values of 0.810, 0.597, and 0.569 for worker retention, tax-financed redistribution, and no government intervention, respectively). The average Infobae effect is likewise indistinguishable from the

Table 5: Minimum detectable effects and equivalence tests for treatment effects

Outcome	Treatment arm	$\hat{\beta}$	SE	MDE (80%)	90% CI
Worker retention	Alarmist	0.010	0.019	0.055	$[-0.022, 0.042]$
	Delayed Threat	0.020	0.019	0.054	$[-0.012, 0.052]$
	Job Reallocation	0.006	0.019	0.054	$[-0.026, 0.037]$
Tax-financed redistribution	Alarmist	0.002	0.010	0.029	$[-0.015, 0.020]$
	Delayed Threat	-0.009	0.010	0.028	$[-0.025, 0.008]$
	Job Reallocation	-0.003	0.010	0.029	$[-0.020, 0.014]$
No government intervention	Alarmist	-0.012	0.018	0.050	$[-0.042, 0.017]$
	Delayed Threat	-0.011	0.018	0.050	$[-0.041, 0.018]$
	Job Reallocation	-0.002	0.018	0.050	$[-0.032, 0.027]$

Notes: $\hat{\beta}$ and SE are the OLS coefficients and robust standard errors from Table 4. The minimum detectable effect (MDE) is computed at 80% power and 5% two-sided significance as $2.80 \times \text{SE}$. The 90% confidence interval is reported to enable two one-sided equivalence tests (TOST) at the 5% level: equivalence to zero within threshold Δ is rejected if the 90% CI is contained in $[-\Delta, +\Delta]$. All nine treatment-outcome combinations reject equivalence at $\Delta = 0.13$ (the effect size reported by Duval et al. (2024) for informational EPL treatments); eight of nine also reject at $\Delta = 0.05$. Reference categories: Control arm and Argentina.

El País effect within each outcome (between-source p -values of 0.475, 0.340, and 0.807). Framing variation does not produce detectable effects even holding source constant, and source variation does not produce detectable effects even averaging across frames.

The ITT null admits two interpretations that our design cannot fully distinguish. The treatments do shift beliefs about the scale or timing of AI displacement, but policy preferences on employment protection are anchored in deeper priors about the state’s redistributive capacity, occupational risk, and identity-linked expectations that new information does not move. Under a competing reading, a single newspaper headline is too thin an informational treatment to shift beliefs at all: the null on preferences would then reflect a weak first stage rather than robust preference stability. Our survey does not include a direct post-treatment manipulation check on the AI framings. A within-survey wording experiment on employment protection regulation, presented in Section 6, provides a bound on the weak-first-stage concern: the same respondents shift their preferences substantially in response to a wording variation of a conceptually adjacent policy question, which is difficult to reconcile with an interpretation in which respondents were uniformly inattentive to the survey.

Pre-treatment beliefs, by contrast, are strongly predictive of policy preferences. Respondents who perceive a higher personal risk of job displacement are more likely to support government intervention, both worker-retention regulation and taxation-based redistribution, and significantly less likely to prefer no government intervention. In addition, beliefs about the incidence of displacement across social groups shape the preferred policy instrument: respondents who identify higher-income groups as particularly vulnerable are less supportive of worker-retention mandates and more supportive of redistribution through taxation, relative to those who view lower-income groups as most vulnerable.

Country-specific results (Table B.4 in Online Appendix B) indicate that average framing effects are generally modest, but they also suggest heterogeneity in how particular narratives map into policy choices. In Mexico, the Job Reallocation frame is associated with higher support for worker-retention regulation and lower support for no government intervention, whereas in Brazil the same frame is associated with an increased propensity to prefer no government intervention. Estimates for Argentina and Colombia remain close to zero across policy options. These cross-country differences are consistent with the idea that the salience of the same AI-employment narrative may depend on baseline priors and the national policy environment. The covariate coefficients, while not interpreted causally, reveal stable correlates of policy preferences: women and respondents with higher perceived displacement risk tend to favor intervention, while more right-leaning respondents are relatively more supportive of no government intervention.

The estimates are robust to specification. Estimating the same model as a multinomial logit (Online Appendix C Table C.1) and inspecting average marginal effects (Online Appendix C Table C.2) yields treatment coefficients that remain small and statistically indistinguishable from zero across both intervention outcomes, while preserving the strong associations between pre-treatment beliefs and policy preferences observed in the OLS estimates. Predicted probabilities by experimental arm (Online Appendix Figure C.1) confirm that the modal preference for worker-retention regulation is essentially constant across treatment conditions.

5.2 Heterogeneous effects

Average treatment effects may mask offsetting responses across subgroups. We estimate interacted linear probability models (Equation (2)) for six pre-specified moderators: education, employment status, age, perceived AI vulnerability, political ideology, and gender. The first three proxy for socioeconomic and labor-market position, and therefore for expected adjustment costs from AI displacement; perceived AI vulnerability captures beliefs about which social group bears the incidence of displacement; political ideology conditions baseline preferences over the role of the state; gender captures the substantial gender gradient documented at baseline. The institutional-credibility interpretation we develop in the conclusion has a specific implication for this analysis: subgroups with greater vulnerability should be more responsive to framing, but the movement should run toward protection rather than redistribution, consistent with anchored discount rates rather than updated risk perceptions. We report model-implied predicted probabilities and nominal p -values throughout; given the number of simultaneous comparisons across moderators, outcomes, and treatment arms (approximately $6 \times 3 \times 3 = 54$ tests), individual significance stars should be interpreted with caution.

Three patterns emerge from the interacted estimates (full coefficient estimates in Online Appendix Tables D.1–D.6; predicted probabilities in Online Appendix Figures D.1–D.4). First, framing ef-

fects are concentrated within specific subgroups: respondents with less than high school education, the currently unemployed, mid-career adults (ages 30–44), and those who self-identify as politically Left. Within these groups, the Delayed Threat and Job Reallocation framings increase predicted support for worker-retention regulation by roughly 6–12 percentage points and reduce support for no government intervention by comparable margins. For respondents with more than high school education, the currently employed, and respondents in other age and ideological categories, predicted probabilities are stable across arms and not statistically distinguishable from zero.

Second, none of the six moderators produces a meaningful shift toward tax-financed redistribution. This is the most consequential pattern for our interpretation: even subgroups with the highest expected adjustment costs from AI displacement, including the unemployed and respondents with less than high school education, do not redirect their demand toward redistribution under any framing condition. Predicted probabilities for the redistributive option remain uniformly low across all subgroups and all framing arms. Among Left respondents, who show the strongest baseline support for state intervention and the largest treatment-induced movement, the movement runs toward worker-retention regulation rather than toward redistribution. Among Left respondents, who show the strongest baseline support for state intervention and the largest treatment-induced movement, the movement runs toward worker-retention regulation rather than toward redistribution. This pattern is noteworthy: even respondents whose ideological priors favor redistribution (the Left) do not overcome the institutional-credibility constraint, suggesting that their skepticism about government capacity to deliver transfers transcends ideological identity. This null holds uniformly across all six moderators: whether they proxy for socioeconomic vulnerability (education, employment, age), beliefs about distributional incidence (perceived AI vulnerability), ideological priors over the state’s role (political ideology), or gender.

Third, the heterogeneity pattern speaks to the design concern about manipulation we flagged in Section 2. If respondents largely ignored the framing content, treatment effects should be indistinguishable from zero across all subgroups, not concentrated within particular ones. The selective responsiveness we observe is more consistent with respondents who processed the headlines but interpreted them through priors about the state’s appropriate role than with uniform inattention to the treatment material.

6 Wording Responsiveness: A Within-Survey Test

The interpretation of the null on the AI framings turns on a question our design cannot directly answer: whether respondents processed the informational content of the treatments. As discussed in Section 5, the observed null on policy preferences is consistent both with genuinely stable priors about the appropriate policy response and with a first stage too weak to move beliefs about AI at all. Our survey did not include a post-treatment manipulation check that would settle

this question directly. A second, independent experimental component of the same instrument, however, provides a within-subject test that speaks to the concern.

The survey administered a wording experiment on employment protection regulation. All respondents ($N = 3,995$) were randomly assigned with equal probability to one of three question wordings, each of which asked whether the respondent considered it necessary that firms be restricted in their ability to dismiss workers when economic conditions deteriorate, even if such restrictions might reduce firms' incentives to hire in the future.⁷ The three treatment arms held the substantive policy content of the question constant but varied the wording along two dimensions: the regulatory tool (either a direct predicate, "make it difficult", or a negated predicate, "make it not easy", to dismiss workers) and the policy objective (either a direct predicate, "to prevent", or a negated predicate, "so as not to favor", firms' practice of dismissing workers when economic conditions deteriorate). The three arms were *Direct Wording*, combining direct predicates on both elements; *Partial Negation*, with a negated tool and a direct objective; and *Full Negation*, with negated predicates on both elements. Only three of the four possible combinations of tool and objective wordings were fielded, which permits two clean pairwise contrasts: *Direct Wording* versus *Partial Negation* isolates the tool-wording change (holding the objective form fixed), and *Partial Negation* versus *Full Negation* isolates the objective-wording change (holding the tool form fixed). Assignment is balanced on the same pre-treatment covariates used in the AI analysis (Online Appendix E Table E.1).

For each of the two contrasts, we estimate treatment effects on two outcome measures: the four-point Likert response scale (1 = "not necessary" to 4 = "very necessary") and a binary indicator equal to one if the respondent reports that regulation is "somewhat" or "very" necessary. Both specifications include the same set of pre-treatment covariates and country fixed effects as the AI regressions in Section 5. Ordered logit estimates on the four-point scale, reported in Online Appendix E Table E.2, yield qualitatively identical results.

Table 6 reports the estimates.⁸ The contrast between *Direct Wording* (make it difficult) and *Partial Negation* (make it not easy), which changes the phrasing of the regulatory tool while holding the policy objective fixed, yields an increase in support for regulation of 0.183 points on the Likert scale ($p < 0.001$) and 7.7 percentage points in the share reporting that regulation is "somewhat" or "very" necessary ($p < 0.001$). The contrast between *Partial Negation* and *Full Negation*, which changes the phrasing of the policy objective while holding the tool phrasing fixed, produces no discernible shift: 0.022 points on the Likert scale ($p = 0.56$) and 1.5 percentage points on the binary outcome ($p = 0.42$). The two clean contrasts are estimated on samples of similar size and with comparable standard errors, so the disparity in magnitude reflects the substantive content of

⁷The full text of each wording is reproduced in Online Appendix A.2.

⁸Full coefficient estimates including all pre-treatment controls are reported in Online Appendix Table E.3.

the wording variation rather than differences in statistical precision.

The 7.7-percentage-point effect of the tool-wording contrast, obtained on the same respondents and with comparable standard errors to the AI regressions, exceeds the minimum detectable effect of any AI framing on worker-retention regulation (5.4 percentage points at 80 percent power; Table 5). That the same instrument elicits substantial preference movement on a conceptually adjacent policy question, but not on any of the AI framings, is difficult to reconcile with an interpretation in which respondents were uniformly inattentive to the survey.

Table 6: Wording effects on support for employment protection regulation (OLS estimates)

	<i>Support for employment protection regulation</i>			
	Direct vs. Partial Negation		Partial vs. Full Negation	
	Likert (1–4) (1)	Binary (0/1) (2)	Likert (1–4) (3)	Binary (0/1) (4)
Partial Negation	0.183*** (0.037)	0.077*** (0.019)		
Full Negation			0.022 (0.037)	0.015 (0.019)
Constant	2.510*** (0.232)	0.576*** (0.120)	2.446*** (0.233)	0.503*** (0.116)
Observations	2,634	2,634	2,634	2,634
R-squared	0.034	0.026	0.022	0.019

Notes: Each column reports OLS coefficients and robust standard errors in parentheses. The dependent variable in Columns (1) and (3) is the four-point Likert measure of the necessity of employment protection regulation, coded from 1 (“not necessary”) to 4 (“very necessary”). The dependent variable in Columns (2) and (4) is a binary indicator equal to one if the respondent reports that regulation is “somewhat” or “very” necessary. Columns (1) and (2) restrict the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms. Columns (3) and (4) restrict the sample to the *Partial Negation* and *Full Negation* arms. All specifications include country fixed effects and the full set of pre-treatment individual controls. Controls include gender, age and age squared, education indicators, wealth index, employment status, political ideology, collective-action measures, perceived police-response efficiency, patience, risk tolerance, basic economic-knowledge measures, a news-following indicator, personal displacement risk, and AI vulnerability group. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Two caveats bear stating. First, our design does not permit a direct manipulation check on the AI framings themselves. The wording experiment rules out an interpretation in which the sample as a whole is uniformly inattentive, but it does not rule out one in which respondents attended to the AI headlines while failing to update on the specific dimensions the treatments varied (magnitude, timing, task composition). Second, the wording experiment varies the linguistic form of the outcome question rather than the informational content of a prior stimulus. The two are not equivalent, but they share the common requirement that respondents read and process semantic content before answering. In this sense, the wording experiment sharpens rather than resolves the interpretation of the AI null by making the weak-first-stage reading substantially less plausible without excluding it entirely. Section E.1 in the Online Appendix discusses the heterogeneous results of this experiment and provides further evidence that the results are not easy to reconcile

with an interpretation in which respondents were uniformly inattentive to the survey.

7 Conclusion

In four Latin American democracies, nearly three-quarters of adults prefer employment-preserving regulation over tax-financed redistribution or non-intervention when asked how governments should respond to AI-driven labor displacement. Preferences vary by country, ideology, and perceptions of AI job risk exposure. This preference does not move when respondents receive information ranging from alarmist forecasts of mass job loss to reassuring accounts of gradual task reallocation. Citizen demand concentrates on a single regulatory margin, preserving existing employment relationships, while the broader portfolio, prominent in expert AI-governance debates (redistribution, retraining support, competition policy), attracts much weaker public support.

The experimental null is itself a substantive finding for AI governance. Because the framing treatments span a wide range of credible narratives about AI’s labor-market impact, the absence of average movement is inconsistent with the view that the gap between citizen demand and the policy mix favored in expert debates reflects specific, correctable misbeliefs about the technology. Whether respondents believe AI threatens up to 60 percent of jobs, that major effects are a decade away, or that the technology will mainly remove low-value work, they converge on employment protection as the preferred policy instrument. The preference appears to be anchored in something more durable than information about the scale or timing of disruption.

The heterogeneity pattern reinforces this interpretation. Where framing has any effect, it operates among the more vulnerable subgroups in the sample: those with less than a high school education, the currently unemployed, mid-career adults, and respondents on the political Left. Yet within these groups, movement runs from non-intervention toward worker-retention regulation, not toward redistribution. Subgroups facing the highest expected adjustment costs from AI displacement do not respond to risk-salient framings by demanding more compensation through taxation; they respond by demanding more protection of existing jobs. Demand for governance instruments appears stable across the dimensions along which one might expect substitution to occur.

The interpretation of this null does not rest on the AI framings alone. A within-survey wording experiment on employment protection regulation, administered on the same instrument and the same respondents, elicits a preference shift of 7.7 percentage points, larger than any AI-framing effect detected here. The sensitivity is selective, responding to the wording of the regulatory tool but not to the wording of the policy objective, and it concentrates in a subset of respondents (employed, middle-aged, non-Left) that is partly disjoint from those most responsive to the AI framings. Both features are difficult to reconcile with a uniform-inattention interpretation of the AI null: preference movement is possible on this instrument; it does not occur in response to

informational variation about AI-driven labor displacement.

One plausible interpretation of these results involves institutional credibility, a question long central to the literature on regulatory governance and credible commitment (Capelos et al., 2016; Levy & Spiller, 1994, 1996; Verhoest et al., 2025). If citizens doubt that their government can tax efficiently and redistribute reliably, demanding visible employment protection is consistent with a rational response to AI risk, even when the welfare logic points the other way. Tax-financed redistribution requires citizens to trade an immediate cost (taxation of those who remain employed) for a future benefit (transfers to those displaced), with the present value of the benefit discounted by the perceived probability that the government will actually deliver. When that probability is low, the implied discount factor approaches zero, and protective regulation that delivers insurance through visible employment guarantees becomes the more credible instrument. Latin America’s persistently low levels of institutional trust (Keefer & Scartascini, 2022; Keefer et al., 2022) are consistent with this interpretation, as is the near-absence of demand for tax-financed redistribution we observe even among subgroups whose ideological priors would otherwise favor it. We cannot test this mechanism directly with our cross-sectional data, but the pattern of results is consistent with it, and we treat it as the most plausible direction for follow-up work. A natural design would experimentally vary the credibility of the redistributive commitment itself by, for example, varying the institutional architecture of the proposed transfer (constitutional protection, autonomous administrator, earmarked revenue source), rather than, as we do here, the framing of the underlying risk.

Two practical implications follow for AI governance design in middle-income democracies. The first is that the gap between politically feasible and welfare-favored AI policy is unlikely to close through communication campaigns or better public information. Closing it would require either building the institutional credibility that makes redistributive commitments believable, or designing instruments that deliver insurance without relying on the state’s discretionary redistributive capacity. Neither is straightforward, but both are more tractable than shifting entrenched preferences through messaging.

The second is more constructive. If citizens prefer protection because they discount redistributive promises, then governance instruments that combine protective features with redistribution-like outcomes may have higher political feasibility than pure redistribution would. Several instruments in active policy debate share this property: mandatory advance-notice and severance schemes with public backstops; sectoral adjustment funds financed by automation-related taxes but disbursed through visible enrollment criteria rather than discretionary transfer; portable benefits and skills accounts that follow the worker rather than the employer; and ex ante regulatory commitments, encoded in legislation or, where institutionally possible, in constitutional or supranational form, that pre-commit to specific compensation structures and thereby reduce the credibility discount

citizens apply to future transfers. None is a perfect substitute for unconditional fiscal redistribution, but each is structured to deliver insurance through a more visible and less discretionary channel. Whether instruments of this kind can attract the public support that pure redistribution lacks is, in our view, the most policy-relevant question raised by our results.

Several limitations bound these conclusions. Our sample is online and skewed toward more-educated urban respondents, which likely understates demand for protection in the broader electorate rather than overstating it, but the magnitude is unknown. The treatments use real newspaper headlines and capture bundled exposure rather than isolated framing effects; we did not administer a post-treatment manipulation check on the AI framings, and while the within-survey wording experiment discussed above provides an indirect bound on the resulting concern, it does not substitute for a direct manipulation check on the AI stimuli themselves. The credibility mechanism is consistent with our data but not directly tested. Each of these points to a productive next step. A longer-form experimental design that varies the institutional architecture of the redistributive commitment, combined with individual-level measures of institutional trust, would identify the credibility mechanism more directly than the framing approach used here. Equally valuable would be evidence from democracies with higher state capacity but similar AI exposure, which would clarify whether the pattern we document is specific to weak-credibility settings or reflects a more general feature of public preferences over AI governance. These caveats notwithstanding, the broader implication for AI governance scholarship seems robust: the political feasibility of welfare-favored AI policy depends not only on the technology and its labor-market effects, but on whether the state can credibly commit to the redistributive promises that such policy entails. In low-credibility settings, governance regimes that pre-commit to specific compensation structures, or that deliver insurance through protection rather than transfer, are likely to attract broader public support than do instruments that require sustained redistributive capacity. Building that capacity remains a long-term project. In the interim, designing AI governance around the credibility constraints citizens actually impose is, we suggest, both more honest and more politically productive than is communicating around them.

8 References

- Acemoglu, D., Autor, D., & Johnson, S. (2026). *Building pro-worker artificial intelligence* (Working Paper No. 34854). National Bureau of Economic Research. <https://doi.org/10.3386/w34854>
- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542. <https://doi.org/10.1257/aer.20160696>
- Athey, S., & Scott Morton, F. (2026, September). Artificial intelligence, competition, and welfare. In *The economics of transformative AI* (pp. 95–114). University of Chicago Press.
- Benish, A., & Levi-Faur, D. (2020). The expansion of regulation in welfare governance. *The ANNALS of the American Academy of Political and Social Science*, 691(1), 17–29. <https://doi.org/10.1177/0002716220949230>
- Bonvecchi, A., Mejia, G., & Scartascini, C. (2026). *Who pays, not when: Fiscal illusion and pension reform preferences in Latin America* (Mimeo). Inter-American Development Bank.
- Braithwaite, J. (2006). Responsive regulation and developing economies. *World Development*, 34(5), 884–898. <https://doi.org/10.1016/j.worlddev.2005.04.021>
- Cafferata, F. G., Gingerich, D. W., & Scartascini, C. (2023). Tempering the taste for vengeance: Information about prisoners and policy choices in Chile. *Comparative Political Studies*, 56(10), 1506–1536. <https://doi.org/10.1177/00104140221141836>
- Capelos, T., Provost, C., Parouti, M., Barnett, J., Chenoweth, J., Fife-Schaw, C., & Kelay, T. (2016). Ingredients of institutional reputations and citizen engagement with regulators. *Regulation & Governance*, 10(4), 350–367. <https://doi.org/10.1111/rego.12097>
- Domínguez, P., & Scartascini, C. (2024). Willingness to pay for crime reduction: The role of information in the Americas. *Journal of Public Economics*, 239, 105205. <https://doi.org/10.1016/j.jpubeco.2024.105205>
- Dussauge-Laguna, M. I. (2024). Regulation and development: Theoretical contributions and empirical lessons from Latin America. *Regulation & Governance*. <https://doi.org/10.1111/rego.12584>
- Duval, R., Ji, Y., Papageorgiou, C., Shibata, I., & Spilimbergo, A. (2024). Preferences for labour regulation: Endowments versus beliefs. *Economic Policy*, 39(119), 549–606. <https://doi.org/10.1093/epolic/eiae021>
- Executive Department, State of California. (2026, May). Executive Order N-6-26 [Signed by Governor Gavin Newsom]. Retrieved May 22, 2026, from <https://www.gov.ca.gov/wp-content/uploads/2026/05/5.21.26-AI-Workforce-EO-FINAL-SIGNED.pdf>
- Gallego, A., Kuo, A., Manzano, D., & Fernández-Albertos, J. (2022). Technological risk and policy preferences. *Comparative Political Studies*, 55(1), 60–92. <https://doi.org/10.1177/00104140211024290>

- Gingerich, D. W., & Scartascini, C. (2022). A heavy hand or a helping hand? Information provision and citizen preferences for anti-crime policies. *Journal of Public Policy*, 42(2), 364–389. <https://doi.org/10.1017/S0143814X21000246>
- Golin, M., & Rauh, C. (2025). *The impact of fear of automation* (Working Paper). <https://doi.org/10.2139/ssrn.5149312>
- Guerreiro, J., Rebelo, S., & Teles, P. (2022). Should robots be taxed? *The Review of Economic Studies*, 89(1), 279–311. <https://doi.org/10.1093/restud/rdab019>
- Horton, D. P., & Lynch-Wood, G. (2024). A resource-based perspective on the regulatory welfare state: Social security in the United Kingdom. *Regulation & Governance*, 18(3), 953–969. <https://doi.org/10.1111/rego.12559>
- Jeffrey, K. (2021). Automation and the future of work: How rhetoric shapes the response in policy preferences. *Journal of Economic Behavior & Organization*, 192, 417–433. <https://doi.org/10.1016/j.jebo.2021.10.019>
- Jeffrey, K., & Matakos, K. (2024). Automation anxiety, fairness perceptions, and redistribution: Past experiences condition the response to future job loss. *Journal of Economic Behavior & Organization*, 221, 174–190. <https://doi.org/10.1016/j.jebo.2024.02.038>
- Keefer, P., & Scartascini, C. (2022). *Trust: The key to social cohesion and growth in Latin America and the Caribbean*. Inter-American Development Bank. <https://doi.org/10.18235/0003792>
- Keefer, P., Scartascini, C., & Vlaicu, R. (2022). Demand-side determinants of public spending allocations: Voter trust, risk and time preferences. *Journal of Public Economics*, 206, 104579. <https://doi.org/10.1016/j.jpubeco.2021.104579>
- Korinek, A., & Lockwood, B. (2026). *Public finance in the age of AI: A primer* (Working Paper No. 34873). National Bureau of Economic Research. <https://doi.org/10.3386/w34873>
- Korinek, A., & Stiglitz, J. E. (2025). *Steering technological progress* (Working Paper No. 232). Institute for New Economic Thinking. <https://doi.org/10.36687/inetwp232>
- Laux, J., Wachter, S., & Mittelstadt, B. (2024). Trustworthy artificial intelligence and the European Union AI act: On the conflation of trustworthiness and acceptability of risk. *Regulation & Governance*, 18(1), 3–32. <https://doi.org/10.1111/rego.12512>
- Lee, J. (2024). Luddite or technophile?—policy preferences for governing technology-driven economic change. *Socio-Economic Review*, 22(3), 1019–1046. <https://doi.org/10.1093/ser/mwae025>
- Levy, B., & Spiller, P. T. (1994). The institutional foundations of regulatory commitment: A comparative analysis of telecommunications regulation. *Journal of Law, Economics, and Organization*, 10(2), 201–246. <https://doi.org/10.1093/oxfordjournals.jleo.a036849>
- Levy, B., & Spiller, P. T. (Eds.). (1996). *Regulations, institutions, and commitment: Comparative studies of telecommunications*. Cambridge University Press.
- Lundgren, M., & Tallberg, J. (2026). *Governing artificial intelligence: Public preferences and regulatory options* [Preprint]. arXiv: 2607.14585. <https://arxiv.org/abs/2607.14585>

- Magistro, B., Loewen, P., Bonikowski, B., Borwein, S., & Lee-Whiting, B. (2024). Attitudes toward automation and the demand for policies addressing job loss: The effects of information about trade-offs. *Political Science Research and Methods*, 12(4), 783–798. <https://doi.org/10.1017/psrm.2024.1>
- Maman, L., Fahy, L., Grimmelhuijsen, S., & Kappler, M. (2025). Measuring citizen trust in regulatory agencies: A systematic review and ways forward. *Regulation & Governance*, 19. <https://doi.org/10.1111/rego.12618>
- Minogue, M., & Cariño, L. (2006). Introduction: Regulatory governance in developing countries. In M. Minogue & L. Cariño (Eds.), *Regulatory governance in developing countries* (pp. 3–16). Edward Elgar Publishing.
- PRISA. (2025, January 24). *EL PAÍS exceeds 400,000 subscribers in just four and a half years*. PRISA. Retrieved February 27, 2026, from <https://www.prisa.com/el-pais-exceeds-400000-subscribers-in-just-four-and-a-half-years/>
- Siddiqui, F. (2026). AI leaders see mass job loss coming. they want government’s help solving it. *The Washington Post*. Retrieved May 22, 2026, from <https://www.washingtonpost.com/technology/2026/05/22/elon-musks-ai-utopia-depends-massive-government-checks/>
- Svirnovskiy, G. (2026). Bill Gates warns AI could lead to mass unemployment, cyberattacks. *Politico*. Retrieved August 26, 2026, from <https://www.politico.com/news/2026/08/26/bill-gates-ai-warning>
- U.S. House of Representatives. (2026, August). H. R. 10044. Retrieved August 6, 2026, from <https://www.congress.gov/bill/119th-congress/house-bill/10044/text>
- van Hoorn, A. (2022). Automatability of work and preferences for redistribution. *Oxford Bulletin of Economics and Statistics*, 84(1), 130–157. <https://doi.org/10.1111/obes.12460>
- Verhoest, K., Maggetti, M., Guaschino, E., & Wynen, J. (2025). How trust matters for the performance and legitimacy of regulatory regimes: The differential impact of watchful trust and good-faith trust. *Regulation & Governance*, 19. <https://doi.org/10.1111/rego.12596>
- WashPostPR. (2020). Arc Publishing drives rapid innovation and 200% growth for Argentina’s Infobae. *The Washington Post*. Retrieved February 27, 2026, from <https://www.washingtonpost.com/pr/2020/01/29/arc-publishing-drives-rapid-innovation-200-growth-argentinas-infobae/>
- Zhang, B., & Dafoe, A. (2020). U.S. public opinion on the governance of artificial intelligence. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES ’20)*, 187–193.

Regulating AI Displacement

Online Appendix

A Survey Questions

A.1 AI-framing experiment

EXPA1A - Group 1 (25% of respondents)

Si la inteligencia artificial realiza tareas que hoy hacen personas y eso hace que se reduzca la cantidad de empleo necesario, Ud que política pública prefiere?

1. Regulaciones que obliguen a contratar trabajadores y conservar empleo
2. Aumentar los impuestos a los que trabajan para subsidiar a los desempleados
3. No introducir regulaciones ni aumentar los impuestos
4. NS/NC

EXPA1A1B - Group 2 (25% of respondents)

Por favor, lea el siguiente titular de un periódico verídico de la región y luego conteste la pregunta.



Si la inteligencia artificial realiza tareas que hoy hacen personas y eso hace que se reduzca la cantidad de empleo necesario, Ud que política pública prefiere?

1. Regulaciones que obliguen a contratar trabajadores y conservar empleo
2. Aumentar los impuestos a los que trabajan para subsidiar a los desempleados
3. No introducir regulaciones ni aumentar los impuestos
4. NS/NC

EXPAIA1C- Group 3 (25% of respondents)

Por favor, lea el siguiente titular de un periódico verídico de la región y luego conteste la pregunta.



EL PAÍS

Economía

MERCADOS · VIVIENDA · FORMACIÓN · MIS DERECHOS · NEGOCIOS · CINCO DÍAS · RETNA · ÚLTIMAS NOTICIAS

INTELIGENCIA ARTIFICIAL >

Goldman Sachs limita el potencial de la inteligencia artificial y cree que su impacto no se notará hasta dentro de diez años

Un artículo del banco de inversión estima que solo el 5% de todas las tareas son susceptibles de automatizarse y cree que no todas las empresas podrán implementar las nuevas tecnologías por su elevado coste

Si la inteligencia artificial realiza tareas que hoy hacen personas y eso hace que se reduzca la cantidad de empleo necesario, Ud que política pública prefiere?

1. Regulaciones que obliguen a contratar trabajadores y conservar empleo
2. Aumentar los impuestos a los que trabajan para subsidiar a los desempleados
3. No introducir regulaciones ni aumentar los impuestos
4. NS/NC

EXPAIA1D - Group 4 (25% of respondents)

Por favor, lea el siguiente titular de un periódico verídico de la región y luego conteste la pregunta.

infobae

Últimas Noticias · Rusia invade Ucrania · Deportes · Venezuela · Tecnología · ESPN · Qué Puedo Ver · Entretenimiento · EEUU

THE NEW YORK TIMES >

Cómo la inteligencia artificial podría erradicar los trabajos sin sentido

El avance de la tecnología promete automatizar tareas que muchos consideran inútiles. ¿Están las personas listas para aceptar estos cambios en sus empleos?

Por Emma Goldberg

06 Ago, 2024 09:30 a.m. EST



Guardar

Si la inteligencia artificial realiza tareas que hoy hacen personas y eso hace que se reduzca la cantidad de empleo necesario, Ud que política pública prefiere?

1. Regulaciones que obliguen a contratar trabajadores y conservar empleo
2. Aumentar los impuestos a los que trabajan para subsidiar a los desempleados
3. No introducir regulaciones ni aumentar los impuestos
4. NS/NC

A.2 Employment protection experiment

REG1A - Group 1 (33% of respondents)

¿Cree usted que es necesario que sea difícil despedir a los empleados que están en las empresas para evitar que las empresas echen trabajadores cuando las condiciones económicas empeoran aun si eso hiciera que las empresas tengan menos incentivos de contratar gente en el futuro?

1. Muy necesario
2. Algo necesario
3. Poco necesario
4. No es necesario

REG1B - Group 2 (33% of respondents)

¿Cree usted que es necesario que no sea fácil despedir a los empleados que están en las empresas para evitar que las empresas echen trabajadores cuando las condiciones económicas empeoran aun si eso hiciera que las empresas tengan menos incentivos de contratar gente en el futuro?

1. Muy necesario
2. Algo necesario
3. Poco necesario
4. No es necesario

REG1C - Group 3 (33% of respondents)

¿Cree usted que es necesario que no sea fácil despedir a los empleados que están en las empresas para no favorecer que las empresas echen trabajadores cuando las condiciones económicas empeoran aun si eso hiciera que las empresas tengan menos incentivos de contratar gente en el futuro?

1. Muy necesario
2. Algo necesario
3. Poco necesario
4. No es necesario

A.3 Experimental controls

AI1

Ahora, para cambiar el tema... La inteligencia artificial es una tecnología que permite a las máquinas aprender y tomar decisiones por sí solas. Se usa, por ejemplo, para reconocer caras, traducir idiomas o predecir el clima. Piensa que la inteligencia artificial puede llegar a reemplazar su trabajo?

1. Si
2. No
3. NS/NC

AI2

Si la inteligencia artificial reemplaza trabajo, eso aumentará el desempleo dentro de

1. Los más pobres
2. La clase media
3. La clase alta
4. NS/NC

AI3

Si un sistema de IA pudiera reducir la delincuencia en su vecindario, pero requiere monitoreo constante de espacios públicos (como cámaras o sensores), ¿Qué tan cómodo(a) se sentiría con ese intercambio entre mayor seguridad y menor privacidad?

1. Muy cómodo
2. Cómodo
3. Indiferente
4. Incómodo
5. Muy incómodo

PROV1T

¿En qué estado/departamento/provincia de (país) vive usted?

MUNICIPIO1T

¿En qué municipio/comuna/ciudad de (nombre del departamento/provincia) vive usted?

Q1TC

Para fines estadísticos, por favor, ¿me podría confirmar su género?

1. Hombre/masculino
2. Mujer/femenino
3. No se identifica como hombre ni como mujer

Q2

¿Cuál es su edad en años cumplidos? _____ años

L1N

Cambiando de tema, tenemos una escala del 1 a 10 que va de izquierda a derecha, en la que el 1 significa izquierda y el 10 significa derecha. Según el sentido que tengan para usted los términos "izquierda" y "derecha" cuando piensa sobre su punto de vista político, ¿dónde se encontraría usted en esta escala?

1. Izquierda
- 2.
- 3.
- 4.
- 5.
- 6.
- 7.
- 8.
- 9.
10. Derecha

EDRE

¿Cuál es el máximo nivel educativo que usted alcanzó?

0. Ninguna
1. Primaria incompleta
2. Primaria completa
3. Secundaria incompleta
4. Secundaria completa
5. Terciaria o universitaria o superior incompleta
6. Terciaria o universitaria o superior completa

988888. No responde

WEALTH1

Imagina una escalera de 10 peldaños, en la cual en el primer peldaño se encuentran los más pobres y en el décimo peldaño, los más ricos. ¿En qué peldaño te ubicarías tú en esta escalera?

0. Primero
- 1.
- 2.
- 3.
- 4.
- 5.
- 6.
- 7.
- 8.
9. Último

I_O1

¿Actualmente, Usted trabaja?

- 0. Sí
- 1. No

I_O1A

¿Cuál es su ocupación principal?

- 0. Asalariado en empresa privada con contrato
- 1. Asalariado en empresa privada sin contrato
- 2. Asalariado/empleado público (del gobierno o empresa estatal)
- 3. Cuenta propia/Independiente/ Profesional liberal
- 4. Patrón o socio de empresa
- 5. Pasantía (remunerado)

I_O1B

¿Ud. estudia, es jubilado, se dedica a actividades de la casa, o está sin empleo?

- 0. Se dedica a las actividades de la casa/ama de casa
- 1. Es jubilado
- 2. Es estudiante
- 3. Está desempleado actualmente
- 4. Está incapacitado permanentemente para trabajar
- 5. No necesita trabajar, tiene otras fuentes de ingreso
- 6. Otros

G1

Muchas de las veredas y calles de la ciudad están en mal estado. Supongamos que el gobierno va a asignar recursos para su mantenimiento a los barrios donde los vecinos logren reunir por lo menos 500 firmas en una petición. ¿Qué tan probable es que en el barrio donde usted vive se logre reunir las 500 firmas para la petición?

- 1. Muy probable
- 2. Algo probable
- 3. Poco probable
- 4. Nada probable

G2

Si logran juntar en el barrio las 500 firmas necesarias para la petición ¿Qué tan probable es que las autoridades efectivamente solucionen el problema?

- 1. Muy probable
- 2. Algo probable
- 3. Poco probable
- 4. Nada probable

H1

Si Ud, tiene una emergencia en su domicilio para la cual necesita llamar a la policía. ¿En cuanto tiempo calcula ud que la policía arribará?

1. Menos de 15 minutos
2. Entre 15 y 30 minutos
3. Más de 30 minutos

RISK1

En su vida personal, ¿en qué medida está dispuesto a renunciar a algo que puede ser beneficioso para Ud. hoy a cambio de obtener mayores beneficios en el futuro? Para ello utilice la siguiente escala de 1 al 5, donde 1 significa que Ud. no está dispuesto a renunciar a obtener los beneficios hoy y 5 significa que Ud. está totalmente dispuesto a esperar a cambio de beneficios mayores en el futuro.

1. Nada dispuesto
- 2.
- 3.
- 4.
5. Muy dispuesto

RISK2

¿Generalmente, es Usted una persona que está dispuesto a correr riesgos o que evita correr riesgos? Para ello utilice la siguiente escala de 1 al 5, donde 1 significa que Ud. no está dispuesto a correr riesgos y 5 significa que Ud. está totalmente dispuesto a correr riesgos.

1. Nada dispuesto
- 2.
- 3.
- 4.
5. Muy dispuesto

GION

¿Con qué frecuencia sigue las noticias, ya sea en la televisión, la radio, los periódicos o el Internet?

1. Diariamente
2. Algunas veces a la semana
3. Algunas veces al mes
4. Algunas veces al año
5. Nunca

988888. No responde

I_L7A

Cambiando otra vez de tema. Si se instalaran más supermercados y tiendas en su comunidad... ¿Qué cree usted que pasaría con los precios de los productos?

1. Aumentarían
2. Se mantendrían igual
3. Disminuirían

I_L9

Imagine que la gente se queja de que los arriendos están muy caros. Ante esta situación el gobierno pone tope al precio de los arriendos. ¿Qué cree usted que pasará?

1. Menos dueños querrán arrendar sus propiedades y habrá menos propiedades disponibles para arrendar
2. Más dueños querrán arrendar sus propiedades y habrá más propiedades para arrendar, o
3. Habrá la misma cantidad de propiedades para arrendar

A.4 Variable construction

Table A.1 documents how the variables used in the empirical analysis are constructed from the survey items listed.

Table A.1: Variable construction from survey items

Variable in analysis	Source item	Coding
Female	Q1TC	1 if “Mujer/femenino”
Employed	I_O1	1 if “Sí” (currently working); 0 otherwise
Personal displacement risk	AI1	1 if “Sí” (AI could replace respondent’s job); 0 otherwise
Understands market competition	I_L7A	1 if “Disminuirían” (correct: more supermarkets lower prices); 0 otherwise
Understands price ceilings	I_L9	1 if option 1 (correct: fewer properties available); 0 otherwise
Covered	I_O1A	1 if “Asalariado en empresa privada con contrato” (formal private) or “Asalariado/empleado público” (public sector); 0 otherwise.
Reg1 binary outcome	REG1	1 if “Muy necesario” or “Algo necesario”; 0 otherwise
Collective action: neighbors act	G1	Reversed; higher = greater perceived likelihood
State responsiveness	G2	Reversed; higher = greater perceived likelihood of state response
Police response efficiency	H1	Reversed; higher = faster expected arrival
Follows the news	GI0N	Reversed; higher = more frequent
Education (3 categories)	EDRE	Less than high school: EDRE $\in$ {“Ninguna”, “Primaria incompleta”, “Primaria completa”, “Secundaria incompleta”}; Completed high school: EDRE = “Secundaria completa”; More than high school: EDRE $\in$ {“Terciaria/universitaria incompleta”, “Terciaria/universitaria completa”}
Political ideology (5 categories)	L1N	Left: L1N $\in$ {1, 2}; Center-left: {3, 4}; Center: {5, 6}; Center-right: {7, 8}; Right: {9, 10}
Age (4 categories)	Q2	18–29; 30–44; 45–59; 60+


B AI-Framing Experiment — Supplementary Results

Table B.1: Balance across respondent characteristics by treatment

	Means				Difference		
	Control (1)	Alarmist (2)	Delayed Threat (3)	Job Reallocation (4)	(2)-(1)	(3)-(1)	(4)-(1)
Female	0.475 (0.500)	0.497 (0.500)	0.518 (0.500)	0.504 (0.500)	0.023 (0.313)	0.043* (0.054)	0.029 (0.201)
Age	40.104 (14.079)	40.616 (13.861)	40.151 (13.938)	40.389 (14.007)	0.515 (0.409)	0.041 (0.948)	0.282 (0.653)
Less than HS education	0.110 (0.313)	0.120 (0.325)	0.117 (0.321)	0.091 (0.288)	0.010 (0.459)	0.007 (0.593)	-0.018 (0.180)
Complete HS education	0.294 (0.456)	0.271 (0.445)	0.280 (0.449)	0.291 (0.455)	-0.022 (0.255)	-0.013 (0.501)	-0.001 (0.954)
More than HS education	0.596 (0.491)	0.609 (0.488)	0.603 (0.489)	0.617 (0.486)	0.012 (0.561)	0.006 (0.777)	0.019 (0.361)
Wealth index (1-10)	4.689 (1.738)	4.680 (1.783)	4.770 (1.826)	4.657 (1.719)	-0.011 (0.885)	0.078 (0.317)	-0.034 (0.651)
Currently employed	0.653 (0.476)	0.650 (0.477)	0.644 (0.479)	0.691 (0.462)	-0.004 (0.864)	-0.009 (0.660)	0.038* (0.072)
Political ideology (1-10)	5.559 (2.495)	5.493 (2.582)	5.572 (2.499)	5.486 (2.589)	-0.066 (0.562)	0.012 (0.916)	-0.076 (0.502)
AI personal displacement risk	0.448 (0.497)	0.441 (0.497)	0.455 (0.498)	0.449 (0.498)	-0.007 (0.763)	0.007 (0.745)	0.002 (0.922)
AI vulnerable group	1.558 (0.570)	1.540 (0.561)	1.562 (0.576)	1.569 (0.593)	-0.017 (0.479)	0.004 (0.879)	0.010 (0.672)
Comfort with AI surveillance	2.066 (1.052)	2.080 (1.069)	2.125 (1.051)	2.062 (1.060)	0.014 (0.774)	0.059 (0.209)	-0.003 (0.949)
Collective action: neighbors act	3.139 (0.874)	3.134 (0.851)	3.069 (0.895)	3.141 (0.869)	-0.005 (0.904)	-0.070* (0.074)	0.002 (0.965)
State responsiveness	2.344 (0.815)	2.355 (0.802)	2.353 (0.814)	2.333 (0.807)	0.011 (0.767)	0.009 (0.805)	-0.010 (0.779)
Police response efficiency	1.696 (0.705)	1.699 (0.676)	1.695 (0.703)	1.669 (0.683)	0.003 (0.909)	-0.001 (0.969)	-0.028 (0.362)
Patience	3.609 (1.087)	3.555 (1.103)	3.571 (1.092)	3.606 (1.100)	-0.055 (0.260)	-0.039 (0.426)	-0.004 (0.935)
Risk tolerance	2.997 (1.145)	2.934 (1.140)	3.022 (1.144)	2.969 (1.149)	-0.064 (0.205)	0.024 (0.632)	-0.028 (0.582)
Understands market competition	0.311 (0.463)	0.341 (0.474)	0.346 (0.476)	0.367 (0.482)	0.031 (0.137)	0.035* (0.095)	0.056*** (0.008)
Understands price ceilings	0.474 (0.500)	0.496 (0.500)	0.512 (0.500)	0.489 (0.500)	0.023 (0.293)	0.039* (0.080)	0.016 (0.471)
Follows the news	4.120 (1.110)	4.082 (1.130)	4.103 (1.114)	4.115 (1.136)	-0.037 (0.458)	-0.017 (0.731)	-0.007 (0.894)
Observations	1,001	999	1,001	995	2,000	2,002	1,996

Notes: Columns (1)–(4) show mean values and standard deviations (in parentheses) for all pre-treatment covariates across the four experimental groups. The Control group corresponds to the control framing condition, while the remaining columns correspond to the three framing treatments (Alarmist, Delayed Threat, and Job Reallocation). Columns (5)–(7) report pairwise mean differences relative to the Control group and corresponding p-values in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Figure B.1: Distribution of preferred policy responses to AI-driven employment risk




Notes: Bars report the share of respondents selecting each policy option (*worker-retention regulation*, *tax-financed redistribution*, or *no government intervention*) among respondents assigned to the control group. Percentages sum to 100.

Table B.2: Baseline support for tax-financed redistribution by labor-market position (control group)

Group	Mean	SE	95% CI	N
<i>By labor status</i>				
Working	5.96%	0.009	[4.1, 7.8]	654
Unemployed	4.35%	0.021	[0.1, 8.6]	92
Out of labor force	5.88%	0.015	[3.0, 8.8]	255
<i>By formality (workers only)</i>				
Formal (contract or public)	4.81%	0.011	[2.6, 7.0]	374
Informal (no contract)	7.69%	0.037	[0.2, 15.2]	52
Self-employed / owner	7.18%	0.018	[3.6, 10.7]	209


Notes: Shares of respondents in the control group who select *tax-financed redistribution* as their preferred policy response to AI-driven labor displacement, by self-reported labor status and, within workers, by formality of employment. Cells report the group mean, robust standard error, 95% confidence interval, and sample size. The formality categorization uses LAPOP occupational codes: formal covers private-sector employees under contract and public-sector employees; informal covers private-sector employees without contract; self-employed groups own-account workers and small business owners.

Figure B.2: Policy support preferences across countries and respondent characteristics




Notes: Bars report the share of respondents selecting each policy option (*worker-retention regulation*, *tax-financed redistribution*, or *no government intervention*) within each group, restricted to respondents assigned to the control group. Percentages sum to 100 within each bar.

Figure B.3: Perceived personal displacement risk from AI




Notes: Bars report the distribution of perceived personal displacement risk from AI, measured by respondents' self-assessment of whether AI could replace their job. Percentages are computed over the full sample ($N = 3,996$).

Figure B.4: Policy support preferences by perceived personal displacement risk from AI




Notes: Bars report the distribution of respondents' preferred policy response to AI-driven employment disruption (*worker-retention regulation*, *tax-financed redistribution*, or *no government intervention*) by perceived personal displacement risk. The risk measure is based on respondents' beliefs about whether AI could replace their jobs. Percentages sum to 100 within each risk category.

Figure B.5: Perceived incidence of AI-related unemployment across social groups



Notes: Bars report the share of respondents indicating which social group would experience the largest increase in unemployment if AI replaces jobs (poor, middle class, or upper class). Percentages are computed over the full sample.

Figure B.6: Policy support preferences by perceived incidence of AI-related unemployment



Notes: Bars report the distribution of preferred policy responses to AI-related employment disruption (*worker-retention regulation*, *tax-financed redistribution*, or *no government intervention*) by the social group respondents identify as most affected by AI-related unemployment. Percentages sum to 100 within each group.

Table B.3: Treatment effects on policy preferences with alternative control sets (OLS estimates)

	<i>Policy preference</i>					
	Worker retention		Tax-financed redistribution		No government intervention	
	(1)	(2)	(3)	(4)	(5)	(6)
Alarmist	0.013 (0.020)	0.010 (0.019)	0.001 (0.010)	0.002 (0.010)	-0.014 (0.018)	-0.012 (0.018)
Delayed Threat	0.022 (0.020)	0.020 (0.019)	-0.008 (0.010)	-0.009 (0.010)	-0.014 (0.018)	-0.011 (0.018)
Job Reallocation	0.012 (0.020)	0.006 (0.019)	-0.005 (0.010)	-0.003 (0.010)	-0.007 (0.018)	-0.002 (0.018)
Personal displacement risk		0.029** (0.014)		0.020*** (0.007)		-0.049*** (0.013)
AI vulnerable group: Middle		-0.012 (0.015)		0.003 (0.008)		0.009 (0.014)
AI vulnerable group: Upper		-0.091** (0.037)		0.053** (0.024)		0.039 (0.033)
Female		0.104*** (0.014)		-0.023*** (0.008)		-0.081*** (0.013)
Age		-0.003 (0.004)		-0.003 (0.002)		0.005 (0.003)
Age sq		0.000 (0.000)		0.000 (0.000)		-0.000 (0.000)
Complete HS education		0.032 (0.025)		-0.021 (0.015)		-0.011 (0.023)
More than HS education		0.013 (0.024)		-0.022 (0.015)		0.009 (0.022)
Wealth index (1-10)		-0.006 (0.004)		0.001 (0.003)		0.005 (0.004)
Currently employed		0.010 (0.016)		-0.000 (0.009)		-0.009 (0.014)
Political ideology (1-10)		-0.025*** (0.003)		0.003* (0.002)		0.022*** (0.003)
Collective action: neighbors act		0.035*** (0.009)		-0.007* (0.004)		-0.028*** (0.008)
State responsiveness		0.010 (0.009)		0.014*** (0.005)		-0.023*** (0.009)
Police response efficiency		-0.016 (0.011)		0.008 (0.005)		0.009 (0.010)
Patience		0.006 (0.007)		-0.014*** (0.004)		0.007 (0.006)
Risk tolerance		0.002 (0.007)		0.005* (0.003)		-0.007 (0.006)
Understands market competition		-0.004 (0.015)		-0.012 (0.007)		0.016 (0.014)
Understands price ceilings		0.014 (0.014)		-0.028*** (0.007)		0.015 (0.013)
Follows the news		0.020*** (0.007)		-0.005 (0.004)		-0.015** (0.006)
Constant	0.670*** (0.019)	0.615*** (0.089)	0.038*** (0.009)	0.163*** (0.050)	0.292*** (0.018)	0.222*** (0.081)
Controls	No	Yes	No	Yes	No	Yes
Country fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,996	3,957	3,996	3,957	3,996	3,957
R-squared	0.009	0.059	0.005	0.032	0.014	0.068

Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Columns (1)–(2) refer to worker-retention regulation, Columns (3)–(4) to redistribution through higher taxes on workers to subsidize the unemployed, and Columns (5)–(6) to no government intervention. Specifications in Columns (2), (4), and (6) include the full set of pre-treatment covariates. Coefficients on treatment indicators (Alarmist, Delayed Threat, Job Reallocation) capture intent-to-treat effects relative to the Control group (omitted category). Country fixed effects are included in all specifications. *** p<0.01, ** p<0.05, * p<0.1.

Table B.4: Treatment effects on policy preferences by country (OLS estimates)

	<i>Policy preference</i>					
	Worker retention		Tax-financed redistribution		No government intervention	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Argentina						
Alarmist	0.030 (0.042)	0.021 (0.041)	0.001 (0.017)	-0.001 (0.017)	-0.030 (0.041)	-0.020 (0.039)
Delayed Threat	0.040 (0.042)	0.058 (0.041)	-0.004 (0.016)	-0.006 (0.016)	-0.036 (0.041)	-0.052 (0.039)
Job Reallocation	0.028 (0.042)	0.044 (0.040)	0.000 (0.017)	0.002 (0.017)	-0.028 (0.041)	-0.046 (0.038)
Personal displacement risk		0.018 (0.029)		0.016 (0.012)		-0.034 (0.027)
AI vulnerable group: Middle		0.014 (0.031)		0.011 (0.012)		-0.025 (0.030)
AI vulnerable group: Upper		-0.103 (0.066)		0.028 (0.028)		0.074 (0.063)
Female		0.136*** (0.031)		-0.029** (0.013)		-0.107*** (0.030)
Age		0.002 (0.008)		-0.004 (0.003)		0.002 (0.008)
Age sq		-0.000 (0.000)		0.000 (0.000)		0.000 (0.000)
Complete HS education		0.133** (0.059)		-0.089*** (0.034)		-0.044 (0.055)
More than HS education		0.022 (0.055)		-0.069** (0.034)		0.046 (0.051)
Wealth index (1–10)		-0.005 (0.009)		-0.003 (0.003)		0.008 (0.008)
Currently employed		0.025 (0.033)		0.011 (0.013)		-0.036 (0.031)
Political ideology (1–10)		-0.048*** (0.006)		-0.001 (0.003)		0.049*** (0.006)
Collective action: neighbors act		0.047*** (0.017)		-0.010 (0.008)		-0.037** (0.016)
State responsiveness		0.006 (0.020)		0.015 (0.009)		-0.021 (0.019)
Police response efficiency		-0.004 (0.021)		0.002 (0.007)		0.002 (0.020)
Patience		-0.003 (0.014)		-0.015** (0.007)		0.018 (0.013)
Risk tolerance		0.011 (0.014)		0.004 (0.006)		-0.015 (0.013)
Understands market competition		-0.025 (0.031)		-0.018 (0.011)		0.043 (0.030)
Understands price ceilings		-0.045 (0.030)		-0.027** (0.012)		0.072** (0.028)
Follows the news		-0.007 (0.014)		0.000 (0.006)		0.007 (0.014)
Constant	0.657*** (0.030)	0.685*** (0.195)	0.036*** (0.012)	0.277*** (0.101)	0.307*** (0.029)	0.038 (0.181)
Controls	No	Yes	No	Yes	No	Yes
Observations	999	989	999	989	999	989
R-squared	0.001	0.126	0.000	0.056	0.001	0.150
Panel B: Brazil						
Alarmist	-0.009	-0.008	-0.008	-0.012	0.017	0.020

	(0.041)	(0.040)	(0.024)	(0.023)	(0.037)	(0.036)
Delayed Threat	0.022	0.013	-0.003	-0.008	-0.018	-0.005
	(0.040)	(0.040)	(0.024)	(0.024)	(0.035)	(0.035)
Job Reallocation	-0.046	-0.064	-0.007	-0.014	0.054	0.077**
	(0.041)	(0.041)	(0.024)	(0.024)	(0.038)	(0.037)
Personal displacement risk		0.016		0.026		-0.042
		(0.030)		(0.018)		(0.027)
AI vulnerable group: Middle		-0.075**		-0.004		0.079**
		(0.036)		(0.020)		(0.033)
AI vulnerable group: Upper		-0.115		0.148*		-0.033
		(0.086)		(0.076)		(0.067)
Female		0.060*		-0.022		-0.037
		(0.031)		(0.019)		(0.028)
Age		-0.001		-0.002		0.003
		(0.008)		(0.004)		(0.007)
Age sq		0.000		-0.000		-0.000
		(0.000)		(0.000)		(0.000)
Complete HS education		0.004		0.033		-0.037
		(0.045)		(0.024)		(0.042)
More than HS education		0.002		0.024		-0.026
		(0.045)		(0.025)		(0.042)
Wealth index (1–10)		-0.018**		0.006		0.011
		(0.009)		(0.006)		(0.008)
Currently employed		-0.021		0.009		0.012
		(0.032)		(0.020)		(0.029)
Political ideology (1–10)		-0.027***		0.004		0.023***
		(0.005)		(0.003)		(0.005)
Collective action: neighbors act		0.013		-0.001		-0.013
		(0.016)		(0.009)		(0.015)
State responsiveness		0.004		0.041***		-0.045***
		(0.018)		(0.012)		(0.016)
Police response efficiency		-0.018		0.018		0.000
		(0.023)		(0.014)		(0.021)
Patience		0.013		-0.012		-0.001
		(0.014)		(0.008)		(0.012)
Risk tolerance		-0.004		0.009		-0.005
		(0.013)		(0.007)		(0.012)
Understands market competition		-0.022		0.007		0.016
		(0.030)		(0.017)		(0.027)
Understands price ceilings		0.022		-0.025		0.004
		(0.029)		(0.017)		(0.026)
Follows the news		0.034**		-0.014		-0.020
		(0.015)		(0.009)		(0.014)
Constant	0.717***	0.755***	0.080***	0.028	0.203***	0.217
	(0.028)	(0.174)	(0.017)	(0.099)	(0.025)	(0.160)
Controls	No	Yes	No	Yes	No	Yes
Observations	999	985	999	985	999	985
R-squared	0.003	0.066	0.000	0.057	0.004	0.070

Panel C: Colombia

Alarmist	0.003	0.010	0.016	0.016	-0.018	-0.026
	(0.037)	(0.037)	(0.018)	(0.018)	(0.034)	(0.034)
Delayed Threat	-0.025	-0.020	-0.000	0.001	0.026	0.020
	(0.038)	(0.037)	(0.017)	(0.017)	(0.035)	(0.034)
Job Reallocation	-0.046	-0.045	0.032	0.032	0.015	0.013
	(0.038)	(0.038)	(0.020)	(0.020)	(0.035)	(0.035)
Personal displacement risk		0.049*		0.021		-0.070***
		(0.027)		(0.014)		(0.024)
AI vulnerable group: Middle		-0.014		0.013		0.001
		(0.027)		(0.013)		(0.025)
AI vulnerable group: Upper		-0.102		0.082		0.020

		(0.082)		(0.056)		(0.075)
Female		0.064**		0.001		-0.065**
		(0.028)		(0.014)		(0.026)
Age		-0.015**		0.001		0.014**
		(0.007)		(0.004)		(0.007)
Age sq		0.000*		-0.000		-0.000*
		(0.000)		(0.000)		(0.000)
Complete HS education		0.128		-0.068		-0.059
		(0.111)		(0.075)		(0.103)
More than HS education		0.130		-0.091		-0.040
		(0.106)		(0.075)		(0.097)
Wealth index (1–10)		-0.004		-0.000		0.004
		(0.009)		(0.006)		(0.008)
Currently employed		0.005		0.017		-0.022
		(0.032)		(0.017)		(0.028)
Political ideology (1–10)		-0.013**		0.003		0.010*
		(0.006)		(0.004)		(0.005)
Collective action: neighbors act		0.021		-0.010		-0.011
		(0.017)		(0.009)		(0.016)
State responsiveness		0.035*		-0.008		-0.026
		(0.019)		(0.008)		(0.018)
Police response efficiency		-0.024		0.011		0.013
		(0.021)		(0.010)		(0.019)
Patience		0.008		-0.002		-0.006
		(0.015)		(0.008)		(0.014)
Risk tolerance		-0.005		-0.001		0.006
		(0.014)		(0.007)		(0.013)
Understands market competition		0.080***		-0.002		-0.078***
		(0.028)		(0.015)		(0.025)
Understands price ceilings		0.002		-0.023		0.021
		(0.028)		(0.014)		(0.025)
Follows the news		0.015		0.001		-0.016
		(0.012)		(0.005)		(0.011)
Constant	0.786***	0.795***	0.036***	0.125	0.177***	0.080
	(0.026)	(0.200)	(0.012)	(0.111)	(0.024)	(0.186)
Controls	No	Yes	No	Yes	No	Yes
Observations	1,000	993	1,000	993	1,000	993
R-squared	0.002	0.051	0.004	0.025	0.002	0.053

Panel D: Mexico

Alarmist	0.029	0.025	-0.004	0.003	-0.025	-0.028
	(0.039)	(0.038)	(0.024)	(0.024)	(0.034)	(0.034)
Delayed Threat	0.051	0.037	-0.024	-0.017	-0.027	-0.020
	(0.038)	(0.038)	(0.022)	(0.023)	(0.034)	(0.033)
Job Reallocation	0.112***	0.099***	-0.043**	-0.031	-0.069**	-0.069**
	(0.036)	(0.036)	(0.021)	(0.021)	(0.032)	(0.032)
Personal displacement risk		0.038		0.007		-0.044*
		(0.027)		(0.016)		(0.024)
AI vulnerable group: Middle		-0.003		-0.007		0.010
		(0.028)		(0.017)		(0.024)
AI vulnerable group: Upper		-0.025		0.018		0.007
		(0.065)		(0.043)		(0.058)
Female	0.127***			-0.047***		-0.080***
	(0.027)			(0.016)		(0.024)
Age		-0.001		-0.005		0.007
		(0.007)		(0.004)		(0.006)
Age sq		0.000		0.000		-0.000
		(0.000)		(0.000)		(0.000)
Complete HS education		0.004		-0.017		0.012
		(0.037)		(0.023)		(0.032)
More than HS education		0.020		-0.023		0.003

Wealth index (1–10)		(0.037)		(0.024)		(0.032)
		0.001		0.003		-0.004
		(0.008)		(0.006)		(0.007)
Currently employed		0.046		-0.042**		-0.004
		(0.030)		(0.019)		(0.026)
Political ideology (1–10)		-0.009*		0.005		0.004
		(0.005)		(0.004)		(0.004)
Collective action: neighbors act		0.057***		-0.005		-0.052***
		(0.018)		(0.010)		(0.017)
State responsiveness		0.002		0.006		-0.008
		(0.019)		(0.012)		(0.017)
Police response efficiency		-0.013		0.004		0.010
		(0.022)		(0.012)		(0.020)
Patience		0.014		-0.021***		0.007
		(0.013)		(0.008)		(0.011)
Risk tolerance		-0.001		0.011*		-0.011
		(0.012)		(0.006)		(0.011)
Understands market competition		-0.014		-0.029*		0.043
		(0.030)		(0.015)		(0.027)
Understands price ceilings		0.079***		-0.038**		-0.041*
		(0.026)		(0.016)		(0.023)
Follows the news		0.029**		-0.005		-0.024**
		(0.012)		(0.008)		(0.011)
Constant	0.733***	0.322**	0.080***	0.314***	0.187***	0.364***
	(0.028)	(0.157)	(0.017)	(0.104)	(0.025)	(0.131)
Controls	No	Yes	No	Yes	No	Yes
Observations	998	990	998	990	998	990
R-squared	0.010	0.074	0.005	0.058	0.005	0.058

Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations of linear probability models. Panels A–D report estimates separately for Argentina, Brazil, Colombia, and Mexico. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Columns (1)–(2) refer to worker-retention regulation, Columns (3)–(4) to tax-financed redistribution to support displaced workers, and Columns (5)–(6) to no government intervention. Specifications in Columns (2), (4), and (6) include the full set of pre-treatment covariates. Coefficients on treatment indicators (Alarmist, Delayed Threat, Job Reallocation) capture intent-to-treat effects relative to the Control group (omitted category). *** p<0.01, ** p<0.05, * p<0.1.

Table B.5: Treatment effects on policy preferences — covered subsample (OLS estimates)

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	0.015 (0.032)	0.003 (0.016)	-0.019 (0.030)
Delayed Threat	-0.002 (0.033)	0.003 (0.016)	-0.001 (0.030)
Job Reallocation	-0.026 (0.031)	-0.000 (0.015)	0.026 (0.029)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	1,514	1,514	1,514

Notes: Each column reports OLS coefficients and robust standard errors (in parentheses) from a linear probability model of the corresponding policy outcome on the treatment arms. The sample is restricted to respondents identifiable as legally covered by the worker-retention regulation (formal private-sector employees under contract and public-sector employees), which account for 38.1% of the full sample. Coefficients on treatment indicators (Alarmist, Delayed Threat, Job Reallocation) capture intent-to-treat effects relative to the Control group (omitted category). All specifications include country fixed effects and the full set of pre-treatment covariates. *** p<0.01, ** p<0.05, * p<0.1.

Table B.6: Minimum detectable effects and equivalence tests for treatment effects on policy preferences — covered subsample

Outcome	Treatment arm	$\hat{\beta}$	SE	MDE (80%)	90% CI
Worker retention	Alarmist	0.015	0.032	0.089	[-0.037, 0.068]
	Delayed Threat	-0.002	0.033	0.092	[-0.056, 0.052]
	Job Reallocation	-0.026	0.031	0.088	[-0.077, 0.026]
Tax-financed redistribution	Alarmist	0.003	0.016	0.044	[-0.023, 0.030]
	Delayed Threat	0.003	0.016	0.045	[-0.023, 0.030]
	Job Reallocation	-0.000	0.015	0.043	[-0.026, 0.025]
No government intervention	Alarmist	-0.019	0.030	0.083	[-0.068, 0.030]
	Delayed Threat	-0.001	0.030	0.085	[-0.051, 0.049]
	Job Reallocation	0.026	0.029	0.082	[-0.022, 0.074]

Notes: $\hat{\beta}$ and SE are the OLS coefficients and robust standard errors on the covered subsample ($n = 1,514$), estimated from linear probability models with country fixed effects and the full set of pre-treatment covariates. The minimum detectable effect (MDE) is computed at 80% power and 5% two-sided significance as $2.80 \times \text{SE}$. The 90% confidence interval is reported to enable two one-sided equivalence tests (TOST) at the 5% level: equivalence to zero within threshold Δ is rejected if the 90% CI is contained in $[-\Delta, +\Delta]$. All nine treatment-outcome combinations reject equivalence at $\Delta = 0.13$ (the effect size reported by Duval et al. (2024) for informational EPL treatments). Reference categories: Control arm and Argentina.

C AI-Framing Experiment — Multinomial Logit

The multinomial logit model treats $Y_i \in \{1, 2, 3\}$ as an unordered categorical outcome and models the probability of each category as:

$$\Pr(Y_i = j) = \frac{\exp(\alpha_j + \mathbf{T}_i' \boldsymbol{\beta}_j + \mathbf{X}_i' \boldsymbol{\gamma}_j + \lambda_{c(i),j})}{\sum_{k=1}^3 \exp(\alpha_k + \mathbf{T}_i' \boldsymbol{\beta}_k + \mathbf{X}_i' \boldsymbol{\gamma}_k + \lambda_{c(i),k})}, \quad (3)$$

where $\mathbf{T}_i$ is a vector of treatment dummies, $\mathbf{X}_i$ is the covariate set used in the LPM, and $\lambda_{c(i),j}$ are country fixed effects. The control group is the omitted treatment category, and no government intervention is the omitted (base) outcome. Equivalently, in log-odds form relative to the base outcome:

$$\log \left(\frac{\Pr(Y_i = j)}{\Pr(Y_i = 3)} \right) = \alpha_j + \mathbf{T}_i' \boldsymbol{\beta}_j + \mathbf{X}_i' \boldsymbol{\gamma}_j + \lambda_{c(i),j}, \quad j \in \{1, 2\}. \quad (4)$$


Table C.1 reports the multinomial logit estimates, with and without the full set of pre-treatment covariates, using *no government intervention* as the base outcome. Treatment coefficients remain small and statistically indistinguishable from zero across both intervention options in both specifications.

In contrast, pre-treatment beliefs and perceived exposure to AI-related labor-market disruption are strongly associated with policy preferences. The coefficient on personal displacement risk is positive and statistically significant for both *worker-retention regulation* and *tax-financed redistribution*, implying that respondents who perceive greater personal vulnerability are more likely to favor some form of government intervention (rather than no intervention). Beliefs about the incidence of displacement across social groups also map into the preferred policy instrument: relative to respondents who perceive lower-income groups as the most vulnerable (the omitted category), those who identify higher-income groups as especially vulnerable are less supportive of worker-retention mandates but more supportive of redistribution via taxes, consistent with substitution across intervention types rather than a uniform increase in demand for government action.

To convey magnitudes in probability terms, Figure C.1 reports average predicted probabilities by experimental arm computed from the multinomial model. The figure confirms a pronounced baseline preference for worker-retention regulation and shows that predicted probabilities remain tightly clustered across the control and treatment groups, with differences that are not statistically distinguishable from zero. Table C.2 reports the corresponding average marginal effects (AMEs) for the treatment indicators.

Table C.3 presents multinomial logit estimates separately by country. Across Argentina, Brazil, Colombia, and Mexico, the estimated coefficients on the treatment indicators are generally small and statistically indistinguishable from zero. The point estimates nonetheless vary across settings, suggesting that any framing effects are not uniform: in Mexico, the Job Reallocation frame is associated with a higher relative propensity to select *worker-retention regulation* (rather than no intervention), whereas in Brazil the same frame is associated with a lower relative propensity to select *tax-financed redistribution*. For Argentina and Colombia, treatment coefficients remain close to zero across both non-baseline categories. Overall, the country-specific results do not provide robust evidence of within-country framing effects, but they are consistent with the possibility that any effects, if present, are context-dependent and modest in magnitude.

Figure C.1: Predicted probabilities of policy support preference by treatment arm



Notes: Bars show predicted probabilities of each policy support preference from a multinomial logit model, by treatment arm (Control, Alarmist, Delayed Threat, Job Reallocation). Vertical lines denote 95% confidence intervals. Horizontal brackets indicate pairwise differences relative to the control arm; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.1: Treatment effects on policy preferences (multinomial logit estimates)

	<i>Policy preference</i>			
	Worker retention		Tax-financed redistribution	
	(1)	(2)	(3)	(4)
Alarmist	0.086 (0.111)	0.068 (0.116)	0.086 (0.210)	0.121 (0.216)
Delayed Threat	0.097 (0.111)	0.092 (0.116)	-0.081 (0.217)	-0.079 (0.225)
Job Reallocation	0.049 (0.110)	0.017 (0.114)	-0.050 (0.214)	-0.017 (0.221)
Personal displacement risk		0.302*** (0.085)		0.633*** (0.159)
AI vulnerable group: Middle		-0.081 (0.091)		-0.014 (0.174)
AI vulnerable group: Upper		-0.305 (0.195)		0.597* (0.325)
Female		0.556*** (0.087)		0.014 (0.166)
Age		-0.035 (0.022)		-0.041 (0.043)
Age sq		0.000 (0.000)		0.000 (0.001)
Complete HS education		0.100 (0.154)		-0.256 (0.262)
More than HS education		-0.048 (0.145)		-0.395 (0.263)
Wealth index (1–10)		-0.033 (0.025)		-0.006 (0.051)
Currently employed		0.059 (0.095)		0.051 (0.185)
Political ideology (1–10)		-0.147*** (0.016)		-0.070** (0.033)
Collective action: neighbors act		0.187*** (0.049)		0.003 (0.094)
State responsiveness		0.134** (0.056)		0.374*** (0.110)
Police response efficiency		-0.071 (0.064)		0.083 (0.116)
Patience		-0.030 (0.041)		-0.292*** (0.078)
Risk tolerance		0.042 (0.040)		0.151** (0.073)
Understands market competition		-0.076 (0.086)		-0.297* (0.176)
Understands price ceilings		-0.036 (0.082)		-0.583*** (0.161)
Follows the news		0.114*** (0.038)		0.018 (0.066)
Constant	0.821*** (0.098)	1.231** (0.539)	-2.080*** (0.219)	-0.254 (0.985)
Observations	3,996	3,957	3,996	3,957

Notes: Each column reports coefficient estimates from multinomial logit regressions of respondents' preferred policy response to potential AI-driven labor displacement. The dependent variable has three unordered categories: *worker-retention regulation*, *tax-financed redistribution* (higher taxes on employed workers to subsidize displaced workers), and *no government intervention*. *No government intervention* is the omitted (base) outcome. Columns (1) and (3) include treatment indicators and country fixed effects only; Columns (2) and (4) additionally include the full set of pre-treatment covariates. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table C.2: Average Marginal Effects on Policy Support Preference

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
Alarmist	0.008 (0.019)	0.004 (0.010)	-0.011 (0.018)
Delayed Threat	0.020 (0.019)	-0.007 (0.010)	-0.012 (0.018)
Job Reallocation	0.004 (0.019)	-0.002 (0.010)	-0.002 (0.018)
Brazil	0.004 (0.022)	0.041*** (0.011)	-0.045** (0.020)
Colombia	0.075*** (0.020)	0.012 (0.010)	-0.087*** (0.018)
Mexico	0.073*** (0.020)	0.018* (0.010)	-0.091*** (0.019)
Personal displacement risk	0.030** (0.014)	0.020*** (0.007)	-0.050*** (0.013)
AI vulnerable group: Middle	-0.014 (0.015)	0.002 (0.008)	0.012 (0.014)
AI vulnerable group: Upper	-0.090** (0.036)	0.056** (0.025)	0.035 (0.032)
Female	0.101*** (0.014)	-0.021*** (0.007)	-0.080*** (0.013)
Age	-0.005 (0.004)	-0.001 (0.002)	0.005 (0.003)
Age sq	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Complete HS education	0.029 (0.025)	-0.019 (0.014)	-0.011 (0.023)
More than HS education	0.008 (0.024)	-0.020 (0.014)	0.012 (0.022)
Wealth index (1–10)	-0.006 (0.004)	0.001 (0.002)	0.005 (0.004)
Currently employed	0.009 (0.016)	0.000 (0.008)	-0.009 (0.015)
Political ideology (1–10)	-0.024*** (0.003)	0.002 (0.001)	0.022*** (0.002)
Collective action: neighbors act	0.034*** (0.008)	-0.007* (0.004)	-0.027*** (0.007)
State responsiveness	0.010 (0.009)	0.013*** (0.005)	-0.023*** (0.008)
Police response efficiency	-0.016 (0.010)	0.007 (0.005)	0.009 (0.010)
Patience	0.006 (0.007)	-0.013*** (0.004)	0.007 (0.006)
Risk tolerance	0.002 (0.007)	0.006* (0.003)	-0.008 (0.006)
Understands market competition	-0.002 (0.014)	-0.012 (0.008)	0.014 (0.013)
Understands price ceilings	0.016 (0.014)	-0.028*** (0.007)	0.011 (0.013)
Follows the news	0.020*** (0.006)	-0.004 (0.003)	-0.017*** (0.006)
Observations	3,957	3,957	3,957

Notes: Each column reports average marginal effects (AMEs) computed from a multinomial logit model of respondents' preferred policy response to potential AI-driven labor displacement. The dependent variable has three unordered categories: *worker-retention regulation* (Column 1), *tax-financed redistribution* (Column 2), and *no government intervention* (Column 3). *No government intervention* is the omitted (base) outcome. The model includes treatment indicators, country fixed effects, and the full set of pre-treatment covariates. AMEs are computed at the observed covariate values and averaged across respondents. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table C.3: Treatment effects on policy preferences by country (multinomial logit estimates)

	<i>Policy preference</i>			
	Worker retention		Tax-financed redistribution	
	(1)	(2)	(3)	(4)
Panel A: Argentina				
Alarmist	0.148 (0.199)	0.123 (0.221)	0.124 (0.500)	0.151 (0.560)
Delayed Threat	0.183 (0.199)	0.337 (0.224)	0.007 (0.514)	0.227 (0.592)
Job Reallocation	0.137 (0.198)	0.283 (0.218)	0.095 (0.500)	0.465 (0.559)
Personal displacement risk		0.182 (0.160)		0.607 (0.387)
AI vulnerable group: Middle		0.144 (0.174)		0.517 (0.454)
AI vulnerable group: Upper		-0.409 (0.326)		0.800 (0.684)
Female		0.647*** (0.169)		-0.399 (0.448)
Age		-0.014 (0.047)		-0.057 (0.114)
Age sq		0.000 (0.001)		-0.000 (0.001)
Complete HS education		0.299 (0.350)		-1.978*** (0.768)
More than HS education		-0.269 (0.322)		-1.628*** (0.570)
Wealth index (1–10)		-0.044 (0.048)		-0.145 (0.105)
Currently employed		0.179 (0.190)		0.554 (0.473)
Political ideology (1–10)		-0.292*** (0.040)		-0.260*** (0.090)
Collective action: neighbors act		0.241*** (0.092)		-0.122 (0.246)
State responsiveness		0.090 (0.108)		0.507* (0.295)
Police response efficiency		-0.027 (0.118)		0.007 (0.250)
Patience		-0.075 (0.079)		-0.505*** (0.179)
Risk tolerance		0.078 (0.079)		0.192 (0.197)
Understands market competition		-0.169 (0.165)		-0.851* (0.461)
Understands price ceilings		-0.376** (0.165)		-1.241*** (0.441)
Follows the news		-0.022 (0.088)		0.063 (0.188)
Constant	0.762*** (0.138)	2.332** (1.116)	-2.147*** (0.352)	3.902 (2.853)
Controls	No	Yes	No	Yes
Observations	999	989	999	989
Panel B: Brazil				
Alarmist	-0.092 (0.221)	-0.112 (0.234)	-0.181 (0.379)	-0.262 (0.398)

Delayed Threat	0.125 (0.229)	0.067 (0.242)	0.052 (0.380)	-0.074 (0.411)
Job Reallocation	-0.302 (0.216)	-0.455** (0.228)	-0.332 (0.375)	-0.508 (0.404)
Personal displacement risk		0.252 (0.176)		0.639** (0.290)
AI vulnerable group: Middle		-0.487*** (0.187)		-0.441 (0.339)
AI vulnerable group: Upper		0.048 (0.562)		1.482** (0.713)
Female		0.276 (0.177)		-0.115 (0.313)
Age		-0.013 (0.045)		-0.011 (0.075)
Age sq		0.000 (0.001)		-0.000 (0.001)
Complete HS education		0.212 (0.262)		0.750 (0.502)
More than HS education		0.147 (0.254)		0.599 (0.507)
Wealth index (1–10)		-0.090* (0.050)		0.013 (0.090)
Currently employed		-0.088 (0.188)		0.099 (0.348)
Political ideology (1–10)		-0.148*** (0.029)		-0.058 (0.052)
Collective action: neighbors act		0.081 (0.094)		0.022 (0.154)
State responsiveness		0.229** (0.107)		0.781*** (0.186)
Police response efficiency		-0.038 (0.132)		0.262 (0.225)
Patience		0.019 (0.081)		-0.166 (0.140)
Risk tolerance		0.030 (0.076)		0.163 (0.125)
Understands market competition		-0.128 (0.170)		0.004 (0.304)
Understands price ceilings		0.027 (0.167)		-0.359 (0.293)
Follows the news		0.146* (0.079)		-0.098 (0.116)
Constant	1.261*** (0.159)	1.277 (1.022)	-0.936*** (0.264)	-2.296 (1.655)
Controls	No	Yes	No	Yes
Observations	999	985	999	985

Panel C: Colombia

Alarmist	0.111 (0.241)	0.165 (0.251)	0.463 (0.486)	0.545 (0.490)
Delayed Threat	-0.168 (0.230)	-0.132 (0.236)	-0.148 (0.515)	-0.100 (0.509)
Job Reallocation	-0.140 (0.233)	-0.148 (0.243)	0.549 (0.462)	0.579 (0.472)
Personal displacement risk		0.477*** (0.176)		0.861*** (0.328)
AI vulnerable group: Middle		-0.024 (0.177)		0.309 (0.353)
AI vulnerable group: Upper		-0.272 (0.483)		1.037 (0.723)
Female		0.457**		0.405

		(0.181)		(0.353)
Age		-0.106**		-0.059
		(0.048)		(0.102)
Age sq		0.001*		0.000
		(0.001)		(0.001)
Complete HS education		0.566		-0.353
		(0.673)		(0.912)
More than HS education		0.406		-1.012
		(0.614)		(0.861)
Wealth index (1–10)		-0.031		-0.019
		(0.052)		(0.134)
Currently employed		0.131		0.499
		(0.202)		(0.473)
Political ideology (1–10)		-0.073**		0.004
		(0.037)		(0.084)
Collective action: neighbors act		0.083		-0.164
		(0.102)		(0.192)
State responsiveness		0.192		-0.042
		(0.126)		(0.208)
Police response efficiency		-0.099		0.158
		(0.134)		(0.242)
Patience		0.037		-0.033
		(0.094)		(0.188)
Risk tolerance		-0.040		-0.036
		(0.090)		(0.185)
Understands market competition		0.579***		0.420
		(0.199)		(0.375)
Understands price ceilings		-0.105		-0.585*
		(0.178)		(0.349)
Follows the news		0.114		0.121
		(0.075)		(0.130)
Constant	1.489***	2.366*	-1.587***	0.127
	(0.167)	(1.297)	(0.366)	(2.310)
Controls	No	Yes	No	Yes
Observations	1,000	993	1,000	993

Panel D: Mexico

Alarmist	0.179	0.178	0.085	0.230
	(0.237)	(0.252)	(0.385)	(0.423)
Delayed Threat	0.224	0.172	-0.195	-0.128
	(0.238)	(0.251)	(0.410)	(0.457)
Job Reallocation	0.601**	0.596**	-0.316	-0.069
	(0.257)	(0.268)	(0.466)	(0.496)
Personal displacement risk		0.333*		0.359
		(0.188)		(0.336)
AI vulnerable group: Middle		-0.086		-0.306
		(0.196)		(0.355)
AI vulnerable group: Upper		-0.096		0.117
		(0.457)		(0.729)
Female		0.737***		-0.277
		(0.197)		(0.355)
Age		-0.041		-0.072
		(0.048)		(0.084)
Age sq		0.000		0.000
		(0.001)		(0.001)
Complete HS education		-0.084		-0.383
		(0.259)		(0.419)
More than HS education		0.008		-0.337
		(0.263)		(0.429)
Wealth index (1–10)		0.033		0.078
		(0.053)		(0.099)

Currently employed	0.089			-0.664*
	(0.217)			(0.367)
Political ideology (1–10)	-0.044			0.063
	(0.037)			(0.074)
Collective action: neighbors act	0.403***			0.250
	(0.122)			(0.211)
State responsiveness	0.048			0.115
	(0.132)			(0.246)
Police response efficiency	-0.085			0.018
	(0.158)			(0.240)
Patience	-0.027			-0.396***
	(0.090)			(0.154)
Risk tolerance	0.079			0.300**
	(0.088)			(0.146)
Understands market competition	-0.273			-0.814*
	(0.200)			(0.428)
Understands price ceilings	0.407**			-0.335
	(0.186)			(0.333)
Follows the news	0.202***			0.105
	(0.075)			(0.132)
Constant	1.365***	-0.247	-0.854***	0.500
	(0.164)	(1.075)	(0.267)	(1.755)
Controls	No	Yes	No	Yes
Observations	998	990	998	990

Notes: Each column reports multinomial logit coefficients and robust standard errors in parentheses. Panels A–D report estimates separately for Argentina, Brazil, Colombia, and Mexico. The dependent variable is respondents’ preferred policy response to potential AI-driven labor displacement, with three categories: *worker-retention regulation*, *tax-financed redistribution*, and *no government intervention*. *No government intervention* is the omitted (baseline) outcome category. Columns (1)–(2) report coefficients for *worker-retention regulation* relative to the baseline, while Columns (3)–(4) report coefficients for *tax-financed redistribution* relative to the baseline. Specifications in Columns (2) and (4) include the full set of pre-treatment covariates. Treatment coefficients (Alarmist, Delayed Threat, Job Reallocation) are reported relative to the Control group (omitted category). *** p<0.01, ** p<0.05, * p<0.1.

D AI-Framing Experiment — Heterogeneous Effects

This appendix reports the detailed heterogeneity analysis summarized in Section 5.2. For each of the six pre-specified moderators—education, employment status, perceived AI vulnerability group, age, political ideology, and gender—we report coefficient estimates from the interacted linear probability model in Equation 2 (Tables D.1–D.6). Model-implied predicted probabilities by treatment arm within each subgroup are reported for education, employment, age, and political ideology (Figures D.1–D.4).

D.1 Education gradients in policy responses

Education shapes both perceived exposure to AI displacement and baseline policy priors. We code it in three categories: less than high school, completed high school, and more than high school. Figure D.1 shows model-implied predicted probabilities by arm within each group; Appendix Table D.1 reports coefficient estimates. Panels (a)–(c) indicate that the pooled null effects mask meaningful heterogeneity by schooling, driven primarily by respondents with less than high school education. Within this subgroup, the Delayed Threat and Job Reallocation framings shift policy support away from *no government intervention* (Panel c) and toward *worker-retention regulation* (Panel a). Consistent with Appendix Table D.1, Delayed Threat increases predicted support for worker-retention regulation by about 12 percentage points and reduces the probability of selecting no intervention by roughly 10 percentage points, while Job Reallocation yields a comparably sized increase in worker-retention support and an even larger decline in no-intervention preferences of about 17 percentage points.

By contrast, Figure D.1 Panel (b) shows no statistically distinguishable treatment effects on *tax-financed redistribution* in any education group. This pattern is consistent with Appendix Table D.1, where the redistribution equation yields small, imprecisely estimated treatment and treatment-by-education coefficients, providing no evidence of a systematic education gradient in responsiveness along this margin. In addition, predicted probabilities for redistribution are uniformly low across arms, so any underlying differences translate into small movements in levels. For respondents with completed high school and with more than high school education, predicted probabilities are stable across experimental arms: treatment and control estimates largely overlap and no differences are statistically distinguishable. The pattern across all education groups points in the same direction: where framing has any effect, it draws respondents away from no intervention and toward worker-retention regulation, not toward redistribution.

D.2 Labor-market attachment: employed vs. unemployed


Employment status shapes perceived economic vulnerability and the salience of job-security concerns. Figure D.2 shows model-implied predicted probabilities for unemployed (omitted category) and employed respondents; Appendix Table D.2 reports coefficient estimates. Consistent with Appendix Table D.2, the only clear pattern of heterogeneity appears for Job Reallocation. Among unemployed respondents, Job Reallocation increases support for *worker-retention regulation* by about 7 percentage points and decreases the probability of choosing *no government intervention* by about 7 percentage points, while effects on *tax-financed redistribution* remain close to zero (Figure D.2, Panel b). For employed respondents, by contrast, predicted probabilities are comparatively stable across arms: treatment and control estimates largely overlap, and the corresponding differences are not statistically distinguishable from zero.

Panel (b) shows no statistically distinguishable treatment effects on *tax-financed redistribution* for either employed or unemployed respondents. Predicted probabilities for redistribution are uniformly low across arms, and treatment effects on that outcome are close to zero for both groups. The experiment therefore finds no evidence that labor-market vulnerability, whether from current unemployment or lower education, redirects demand from employment protection toward redistribution.

D.3 Perceived AI vulnerability

Beliefs about which social group bears the incidence of AI-driven displacement may shape responsiveness to informational framing independently of the socioeconomic and ideological dimensions considered above. Respondents who see AI unemployment as concentrated among the poor, the middle class, or the upper class may weight the framing content differently depending on their perceived distributional stakes. Table D.3 reports coefficient estimates from

Figure D.1: Heterogeneous treatment effects by education



Notes: Bars show predicted probabilities of selecting each policy option by education group and treatment arm from an OLS linear probability model that allows treatment effects to vary by education (less than high school, completed high school, more than high school). Panels (a)–(c) report predicted probabilities for worker-retention regulation, tax-financed redistribution, and no government intervention, respectively. The model includes country fixed effects and the full set of pre-treatment covariates. Vertical lines denote 95% confidence intervals. Brackets indicate pairwise tests relative to the control group within each education category. *n.s.* denotes not statistically significant; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

the interacted linear probability model (Equation (2)), with respondents identifying the poor as most vulnerable as the omitted category. None of the six treatment-by-vulnerability interaction coefficients reaches conventional significance in any of the three outcomes. Perceived distributional incidence does not condition responsiveness to any of the framing treatments, consistent with the broader null.

Table D.1: Heterogeneous treatment effects by education (OLS estimates)

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	0.037 (0.062)	0.056 (0.036)	-0.093 (0.057)
Delayed Threat	0.117** (0.059)	-0.013 (0.029)	-0.105* (0.056)
Job Reallocation	0.110* (0.063)	0.057 (0.040)	-0.168*** (0.055)
<i>Education</i>			
Complete HS education	0.083 (0.052)	0.010 (0.026)	-0.093* (0.049)
More than HS education	0.081 (0.050)	0.006 (0.024)	-0.087* (0.047)
<i>Interactions</i>			
Alarmist × Complete HS	-0.036 (0.072)	-0.060 (0.042)	0.096 (0.066)
Alarmist × More than HS	-0.027 (0.067)	-0.062 (0.038)	0.089 (0.062)
Delayed Threat × Complete HS	-0.098 (0.069)	0.014 (0.036)	0.084 (0.064)
Delayed Threat × More than HS	-0.116* (0.064)	0.001 (0.031)	0.115* (0.061)
Job Reallocation × Complete HS	-0.078 (0.072)	-0.084* (0.044)	0.162** (0.064)
Job Reallocation × More than HS	-0.135** (0.068)	-0.059 (0.042)	0.194*** (0.060)
Constant	0.565*** (0.096)	0.139*** (0.051)	0.296*** (0.088)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	3,957	3,957	3,957
R-squared	0.061	0.035	0.071


Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Treatment indicators (Alarmist, Delayed Threat, Job Reallocation) are defined relative to the Control group (omitted category). Education is coded in three categories: less than high school (omitted), completed high school, and more than high school. All specifications include country fixed effects and the full set of pre-treatment covariates. *** p<0.01, ** p<0.05, * p<0.1.

D.4 Age and adjustment horizons

Age proxies for exposure to automation risk and adjustment costs across the life cycle. We code age in four categories, 18–29 (omitted), 30–44, 45–59, and 60+, and report model-implied predicted probabilities in Figure D.3; Appendix Table D.4 reports coefficient estimates. Framing effects are limited, with one exception concentrated among respondents aged 30–44. In this group, both the Alarmist and Delayed Threat framings significantly increase support for *worker-retention regulation* relative to Control (Panel (a)) by roughly 6–7 percentage points. Consistent with this shift toward regulation, the Alarmist framing also significantly reduces the probability of selecting *no government intervention* (Panel (c)) for ages 30–44 by about 6 percentage points. By contrast, Panel (b) shows little movement in *tax-financed redistribution* across most age groups; the only detectable change is a modest decline for ages 30–44 under Delayed Threat, on the order of 3 percentage points.

Outside the 30–44 category, differences are not statistically distinguishable from zero. Age heterogeneity is therefore limited and concentrated in the mid-career cohort, where framing shifts support toward worker-retention regulation and away from no intervention—again without moving redistribution.

Figure D.2: Heterogeneous treatment effects by employment status



Notes: Bars show predicted probabilities of selecting each policy option by employment status and treatment arm from an OLS linear probability model that allows treatment effects to vary by employment status (unemployed/other: omitted category; employed). Panels (a)–(c) report predicted probabilities for worker-retention regulation, tax-financed redistribution, and no government intervention, respectively. The model includes country fixed effects and the full set of pre-treatment covariates. Vertical lines denote 95% confidence intervals. Brackets indicate pairwise tests relative to the control group within each employment-status category. *n.s.* denotes not statistically significant; ****** $p < 0.01$, ***** $p < 0.05$, ***** $p < 0.1$.

D.5 Ideological predispositions

Political ideology conditions baseline preferences over the state's role and receptiveness to AI narratives. Figure D.4 shows model-implied predicted probabilities across a five-category ideology scale (Left, Center-left, Center, Center-right, Right); Appendix Table D.5 reports coefficient estimates. Baseline ideological differences are pronounced—support for regulation declines and support for no intervention rises from left to right—but treatment-induced

Table D.2: Heterogeneous treatment effects by employment status (OLS estimates)

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	-0.022 (0.033)	0.008 (0.018)	0.015 (0.031)
Delayed Threat	0.039 (0.033)	-0.006 (0.017)	-0.033 (0.030)
Job Reallocation	0.069** (0.033)	-0.001 (0.018)	-0.068** (0.029)
<i>Employment status</i>			
Employed	0.027 (0.030)	0.003 (0.016)	-0.030 (0.027)
<i>Interactions</i>			
Alarmist × Employed	0.050 (0.041)	-0.008 (0.022)	-0.042 (0.038)
Delayed Threat × Employed	-0.030 (0.041)	-0.003 (0.021)	0.033 (0.037)
Job Reallocation × Employed	-0.093** (0.040)	-0.003 (0.022)	0.096*** (0.037)
Constant	0.606*** (0.090)	0.160*** (0.051)	0.233*** (0.082)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	3,957	3,957	3,957
R-squared	0.062	0.032	0.071

Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Treatment indicators (Alarmist, Delayed Threat, Job Reallocation) are defined relative to the Control group (omitted category). Employment status is coded as a binary indicator equal to one if the respondent is employed and zero otherwise (unemployed/otherwise; omitted category). All specifications include interactions between treatment indicators and employment status, country fixed effects, and the full set of pre-treatment covariates. *** p<0.01, ** p<0.05, * p<0.1.

differences are modest and concentrated on the left.

In Panel (a), the main detectable heterogeneity emerges among respondents who self-identify as Left. Within this subgroup, the Job Reallocation framing increases predicted support for *worker-retention regulation* relative to the control arm by roughly 11 percentage points, consistent with Table D.5 (the Job Reallocation coefficient is defined relative to the control group within the omitted ideology category). By contrast, treatment-control differences are statistically indistinguishable from zero for Center-left, Center, Center-right, and Right respondents, indicating that the pro-regulation shift is not evident outside the left.

Panel (c) shows a symmetric pattern for *no government intervention*. Among Left respondents, Job Reallocation reduces predicted support for laissez-faire by about 9 percentage points relative to control (Appendix Table D.5). For other ideology groups, estimates largely overlap across arms. Table D.5 also aligns with this visual pattern: interaction terms imply that the left-specific reduction in no-intervention under Job Reallocation attenuates as ideology moves toward the center and right, yielding near-zero net effects in those groups.

By contrast, Panel (b) shows no statistically distinguishable treatment effects on *tax-financed redistribution* in any ideology category. Redistribution remains a minority preference across the entire ideological spectrum, including among left-leaning respondents who show the strongest baseline support for state intervention. This reinforces the broader pattern in the data: AI-related concern translates into demand for employment protection across ideological groups, with redistribution remaining a secondary instrument even where prior ideological dispositions might favor a more redistributive response.

Table D.3: Heterogeneous treatment effects by perceived AI vulnerability (OLS estimates)

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	-0.008 (0.028)	0.007 (0.016)	0.001 (0.025)
Delayed Threat	-0.005 (0.028)	0.000 (0.015)	0.004 (0.025)
Job Reallocation	0.009 (0.027)	-0.008 (0.015)	-0.001 (0.024)
<i>AI vulnerable group</i>			
Middle class	-0.033 (0.029)	0.007 (0.015)	0.026 (0.027)
Upper class	-0.084 (0.077)	0.056 (0.048)	0.029 (0.069)
<i>Interactions</i>			
Alarmist × Middle	0.037 (0.040)	-0.010 (0.021)	-0.027 (0.036)
Alarmist × Upper	0.006 (0.111)	0.002 (0.074)	-0.008 (0.102)
Delayed Threat × Middle	0.050 (0.040)	-0.015 (0.020)	-0.035 (0.036)
Delayed Threat × Upper	0.011 (0.101)	-0.041 (0.063)	0.030 (0.093)
Job Reallocation × Middle	-0.004 (0.040)	0.008 (0.020)	-0.003 (0.037)
Job Reallocation × Upper	-0.039 (0.102)	0.025 (0.065)	0.014 (0.090)
Constant	0.624*** (0.090)	0.160*** (0.051)	0.216*** (0.081)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	3,957	3,957	3,957
R-squared	0.060	0.033	0.068


Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Treatment indicators (Alarmist, Delayed Threat, Job Reallocation) are defined relative to the Control group (omitted category). Perceived AI vulnerable group is coded in three categories: poor (omitted), middle class, and upper class. All specifications include country fixed effects and the full set of pre-treatment covariates. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D.6 Gender

In our baseline specification (Table 4), women are 10.3 percentage points more likely than men to prefer worker-retention regulation ($p < 0.01$), 2.3 percentage points less likely to prefer tax-financed redistribution ($p < 0.01$), and 8.0 percentage points less likely to prefer no government intervention ($p < 0.01$), conditional on all other covariates. Given the magnitude of this baseline gradient and the marginal imbalance in the female share across treatment arms (Table B.1), we examine whether treatment responsiveness itself varies by gender. Table D.6 reports estimates from a specification that interacts each treatment arm with a female indicator. None of the nine treatment-by-gender interaction coefficients reaches conventional significance in any outcome. Treatment effects are therefore statistically indistinguishable across women and men.

The gender gap in baseline preferences is substantial, but it does not reflect differential responsiveness to information. Both women and men respond to the informational treatments in the same way—that is, they do not respond at all—which indicates that the null result we document is not concealing gender-driven heterogeneity: it reflects equally stable priors held by both groups. This pattern is consistent with the broader interpretation that policy preferences

Figure D.3: Heterogeneous treatment effects by age



Notes: Bars show predicted probabilities of selecting each policy option by age group and treatment arm from an OLS linear probability model that allows treatment effects to vary by age (18-29: omitted category; 30-44, 45-59, 60+). Panels (a)–(c) report predicted probabilities for worker-retention regulation, tax-financed redistribution, and no government intervention, respectively. The model includes country fixed effects and the full set of pre-treatment covariates. Vertical lines denote 95% confidence intervals. Brackets indicate pairwise tests relative to the control group within each age category. *n.s.* denotes not statistically significant; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


on employment protection are anchored in deep, structural differences in how women and men perceive labor-market vulnerability, the incidence of AI displacement, and the appropriate role of the state—priors formed well before any exposure to short-term informational cues and unresponsive to them within either gender. The gender gradient, in other words, is a distributional feature of preferences rather than a differential margin of response.

Table D.4: Heterogeneous treatment effects by age (OLS estimates)

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	-0.025 (0.034)	0.009 (0.022)	0.015 (0.030)
Delayed Threat	-0.035 (0.035)	0.028 (0.023)	0.007 (0.030)
Job Reallocation	0.021 (0.033)	-0.003 (0.021)	-0.018 (0.029)
<i>Age</i>			
30–44	-0.082** (0.035)	0.004 (0.021)	0.079** (0.032)
45–59	-0.028 (0.036)	-0.031 (0.019)	0.058* (0.033)
60+	-0.018 (0.049)	-0.032 (0.023)	0.050 (0.045)
<i>Interactions</i>			
Alarmist × 30–44	0.085* (0.050)	-0.014 (0.030)	-0.071 (0.044)
Alarmist × 45–59	0.042 (0.050)	-0.004 (0.027)	-0.038 (0.046)
Alarmist × 60+	-0.046 (0.070)	-0.012 (0.033)	0.058 (0.065)
Delayed Threat × 30–44	0.106** (0.049)	-0.061** (0.029)	-0.045 (0.044)
Delayed Threat × 45–59	0.065 (0.051)	-0.031 (0.028)	-0.033 (0.047)
Delayed Threat × 60+	0.023 (0.069)	-0.062** (0.029)	0.039 (0.065)
Job Reallocation × 30–44	-0.001 (0.050)	0.000 (0.029)	0.001 (0.044)
Job Reallocation × 45–59	-0.029 (0.049)	-0.002 (0.026)	0.031 (0.045)
Job Reallocation × 60+	-0.059 (0.070)	0.003 (0.032)	0.056 (0.065)
Constant	0.588*** (0.062)	0.098*** (0.034)	0.315*** (0.057)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	3,957	3,957	3,957
R-squared	0.062	0.032	0.071

Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Treatment indicators (Alarmist, Delayed Threat, Job Reallocation) are defined relative to the Control group (omitted category). Age is coded in four categories: 18–29 (omitted), 30–44, 45–59, and 60+. All specifications include interactions between treatment indicators and age categories, country fixed effects, and the full set of pre-treatment covariates. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure D.4: Heterogeneous treatment effects by political ideology



Notes: Bars show predicted probabilities of selecting each policy option by political ideology group and treatment arm from an OLS linear probability model that allows treatment effects to vary by ideology (Left: omitted category; Center-left, Center, Center-right, Right). Panels (a)–(c) report predicted probabilities for worker-retention regulation, tax-financed redistribution, and no government intervention, respectively. The model includes country fixed effects and the full set of pre-treatment covariates. Vertical lines denote 95% confidence intervals. Brackets indicate pairwise tests relative to the control group within each ideology category. *n.s.* denotes not statistically significant; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.5: Heterogeneous treatment effects by political ideology (OLS estimates)

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	0.077 (0.052)	-0.011 (0.032)	-0.066 (0.045)
Delayed Threat	0.079 (0.052)	-0.014 (0.033)	-0.064 (0.044)
Job Reallocation	0.110** (0.050)	-0.022 (0.032)	-0.088** (0.042)
<i>Political ideology</i>			
Center-left	0.029 (0.051)	0.008 (0.032)	-0.036 (0.044)
Center	0.007 (0.046)	-0.024 (0.027)	0.017 (0.040)
Center-right	-0.097* (0.055)	-0.013 (0.029)	0.110** (0.050)
Right	-0.164*** (0.056)	0.028 (0.034)	0.136*** (0.050)
<i>Interactions</i>			
Alarmist × Center-left	-0.07 (0.068)	0.017 (0.044)	0.054 (0.057)
Alarmist × Center	-0.081 (0.061)	0.017 (0.035)	0.065 (0.053)
Alarmist × Center-right	-0.109 (0.076)	0.019 (0.040)	0.090 (0.070)
Alarmist × Right	-0.034 (0.075)	-0.001 (0.044)	0.035 (0.067)
Delayed Threat × Center-left	-0.122* (0.070)	-0.014 (0.042)	0.137** (0.060)
Delayed Threat × Center	-0.066 (0.060)	0.026 (0.037)	0.040 (0.052)
Delayed Threat × Center-right	-0.058 (0.073)	-0.011 (0.038)	0.069 (0.065)
Delayed Threat × Right	-0.015 (0.076)	-0.002 (0.045)	0.018 (0.067)
Job Reallocation × Center-left	-0.174*** (0.067)	0.026 (0.043)	0.148** (0.057)
Job Reallocation × Center	-0.106* (0.059)	0.015 (0.035)	0.091* (0.051)
Job Reallocation × Center-right	-0.136* (0.071)	0.063 (0.041)	0.074 (0.064)
Job Reallocation × Right	-0.046 (0.073)	-0.022 (0.042)	0.068 (0.066)
Constant	0.496*** (0.097)	0.186*** (0.055)	0.318*** (0.087)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	3,957	3,957	3,957
R-squared	0.063	0.038	0.072

Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Treatment indicators (Alarmist, Delayed Threat, Job Reallocation) are defined relative to the Control group (omitted category). Political ideology is coded in five categories: Left (omitted), Center-left, Center, Center-right, and Right. All specifications include country fixed effects and the full set of pre-treatment covariates. *** p<0.01, ** p<0.05, * p<0.1.

Table D.6: Heterogeneous treatment effects by gender (OLS estimates)

	<i>Policy preference</i>		
	Worker retention (1)	Tax-financed redistribution (2)	No government intervention (3)
<i>Treatments</i>			
Alarmist	-0.000 (0.029)	0.001 (0.016)	-0.001 (0.027)
Delayed Threat	0.035 (0.029)	-0.016 (0.015)	-0.019 (0.027)
Job Reallocation	-0.014 (0.029)	0.004 (0.016)	0.010 (0.027)
<i>Gender</i>			
Female	0.095*** (0.028)	-0.023 (0.015)	-0.072*** (0.026)
<i>Interactions</i>			
Alarmist $\times$ Female	0.021 (0.039)	0.002 (0.021)	-0.023 (0.035)
Delayed Threat $\times$ Female	-0.028 (0.039)	0.015 (0.020)	0.014 (0.036)
Job Reallocation $\times$ Female	0.040 (0.039)	-0.014 (0.020)	-0.026 (0.036)
Constant	0.620*** (0.089)	0.162*** (0.050)	0.218*** (0.081)
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	3,957	3,957	3,957
R-squared	0.060	0.033	0.068

Notes: Each column reports regression coefficients and robust standard errors in parentheses from OLS estimations. The dependent variable is an indicator equal to one if the respondent selects the corresponding policy option and zero otherwise. Treatment indicators (Alarmist, Delayed Threat, Job Reallocation) are defined relative to the Control group (omitted category). Female is coded as a binary indicator equal to one if the respondent is female and zero otherwise (male; omitted category). All specifications include interactions between treatment indicators and Female, country fixed effects, and the full set of pre-treatment covariates. *** p<0.01, ** p<0.05, * p<0.1.

E Wording Experiment — Supplementary Results

Table E.1: Balance across respondent characteristics — wording experiment

	Means			Difference	
	Direct Wording (1)	Partial Negation (2)	Full Negation (3)	(2)-(1)	(3)-(2)
Female	0.481 (0.500)	0.497 (0.500)	0.517 (0.500)	0.016 (0.419)	0.020 (0.305)
Age	40.165 (14.025)	40.446 (14.021)	40.319 (13.854)	0.281 (0.604)	-0.134 (0.804)
Less than HS Education	0.117 (0.322)	0.109 (0.312)	0.101 (0.302)	-0.008 (0.534)	-0.008 (0.486)
Complete HS Education	0.275 (0.447)	0.278 (0.448)	0.302 (0.459)	0.003 (0.864)	0.024 (0.159)
More than HS Education	0.607 (0.489)	0.613 (0.487)	0.597 (0.491)	0.005 (0.797)	-0.016 (0.371)
Wealth index (1-10)	4.647 (1.801)	4.760 (1.732)	4.686 (1.767)	0.112* (0.098)	-0.073 (0.274)
Currently employed	0.655 (0.475)	0.673 (0.469)	0.650 (0.477)	0.018 (0.325)	-0.023 (0.205)
Political ideology (1-10)	5.541 (2.574)	5.491 (2.471)	5.545 (2.583)	-0.053 (0.587)	0.054 (0.578)
AI Personal displacement risk	0.435 (0.496)	0.456 (0.498)	0.456 (0.498)	0.022 (0.253)	-0.000 (0.991)
AI vulnerable group:	1.554 (0.573)	1.565 (0.574)	1.552 (0.580)	0.012 (0.571)	-0.014 (0.514)
Collective action: neighbors act	3.121 (0.888)	3.109 (0.876)	3.132 (0.853)	-0.011 (0.747)	0.023 (0.502)
State responsiveness	2.360 (0.829)	2.347 (0.798)	2.334 (0.801)	-0.013 (0.672)	-0.013 (0.681)
Police response efficiency	1.702 (0.705)	1.701 (0.686)	1.666 (0.685)	-0.002 (0.933)	-0.035 (0.177)
Patience	3.579 (1.093)	3.601 (1.102)	3.577 (1.093)	0.022 (0.602)	-0.024 (0.580)
Risk tolerance	2.950 (1.162)	2.983 (1.146)	3.010 (1.125)	0.035 (0.434)	0.026 (0.550)
Understands market competition	0.340 (0.474)	0.343 (0.475)	0.339 (0.473)	0.002 (0.895)	-0.004 (0.809)
Understands price ceilings	0.500 (0.500)	0.484 (0.500)	0.493 (0.500)	-0.016 (0.402)	0.009 (0.644)
Follows the news	4.107 (1.130)	4.109 (1.111)	4.098 (1.128)	0.000 (0.997)	-0.011 (0.800)
Observations	1,337	1,326	1,332	2,663	2,658

Notes: Columns (1)–(3) show mean values and standard deviations (in parentheses) for all pre-treatment covariates across the three arms of the wording experiment described in Section 6. Columns (4) and (5) report pairwise mean differences and corresponding p-values in parentheses for the two contrasts: column (4) reports the difference between the *Partial Negation* and *Direct Wording* arms, and column (5) reports the difference between the *Full Negation* and *Partial Negation* arms. A comparison between the Direct Wording and Full Negation arms would confound both dimensions and is therefore not reported. *** p<0.01, ** p<0.05, * p<0.1.

E.1 Heterogeneous wording effects

The aggregate contrast in Table 6 conceals substantial variation in who responds to the tool-wording change. Interacted linear probability models for the same six pre-specified moderators used in the AI heterogeneity analysis (Online Appendix E, Tables E.4–E.9) locate the effect predominantly among respondents currently in the labor force, in the middle age brackets, outside the ideological Left, and among those who identify middle class as the group most vulnerable to AI unemployment. In each of these subgroups, the shift from Direct Wording to Partial Negation produces a larger increase in support for employment protection regulation than in the corresponding baseline category. Employed respondents increase support preferences by an additional 8.5 percentage points relative to the non-employed baseline ($p < 0.05$); respondents aged 30–59 increase support by 9 to 11 additional percentage points relative to the 18–29 baseline ($p < 0.10$ or better); respondents outside the Left category increase support by 11 to 15 percentage points relative to Left respondents ($p < 0.05$ for Center-left); and respondents identifying middle-class as most vulnerable increase support by an additional 8.3 percentage points relative to those identifying the poor ($p < 0.05$). Interactions with education and gender are not statistically distinguishable from zero.

Two features of this heterogeneity pattern bear on the interpretation of the paper’s central null. The first is that the subgroups activated by the tool-wording contrast differ from those activated by the AI framings. Where the AI treatments concentrated their limited movement among lower-education, unemployed, and left-leaning respondents (Section 5.2), the wording change moves preferences in a partly non-overlapping subset defined by labor-market attachment, life-cycle position, ideological location outside the left, and beliefs about the distributional incidence of AI displacement. The second is that the null on the objective-wording contrast holds on average, and no interaction with moderators achieves conventional significance ($p < 0.05$). Neither feature is easy to reconcile with an interpretation in which respondents were uniformly inattentive to the survey: uniform inattention would produce either no heterogeneity or a heterogeneity pattern common to both experimental components, not two distinct ones alongside a wording contrast whose null is itself uniform across subgroups.

Table E.2: Wording effects on support for employment protection regulation (ordered logit)

	<i>Support for employment protection regulation</i>	
	Direct vs. Partial Negation (1)	Partial vs. Full Negation (2)
Partial Negation	0.360*** (0.072)	
Full Negation		0.058 (0.072)
AI Personal displacement risk	0.021 (0.074)	0.054 (0.075)
AI vulnerable group: Middle class	0.071 (0.080)	0.060 (0.077)
AI vulnerable group: Upper class	-0.355* (0.206)	-0.136 (0.184)
Female	0.036 (0.074)	0.029 (0.076)
Age	-0.017 (0.018)	-0.003 (0.018)
Age sq	0.000 (0.000)	-0.000 (0.000)
Complete HS Education	0.166 (0.135)	0.198 (0.141)
More than HS Education	0.099 (0.130)	0.227* (0.138)
Wealth index (1-10)	0.006 (0.023)	-0.013 (0.024)
Currently employed	0.121 (0.082)	0.185** (0.084)
Political ideology (1-10)	-0.042*** (0.016)	-0.037** (0.016)
Collective action: neighbors act	0.101** (0.046)	0.093** (0.046)
State responsiveness	0.090 (0.055)	0.144*** (0.053)
Police response efficiency	0.080 (0.059)	0.022 (0.060)
Patience	0.008 (0.041)	0.014 (0.042)
Risk tolerance	-0.008 (0.039)	-0.017 (0.041)
Understands market competition	0.041 (0.077)	0.130* (0.078)
Understands price ceilings	-0.022 (0.074)	-0.023 (0.073)
Follows the news	0.039 (0.035)	0.047 (0.035)
Cutpoint 1	-1.414*** (0.456)	-1.179*** (0.455)
Cutpoint 2	-0.067 (0.456)	0.123 (0.456)
Cutpoint 3	1.848*** (0.457)	1.999*** (0.459)
Observations	2,634	2,634

Notes: Each column reports ordered logit coefficients and robust standard errors in parentheses. The dependent variable is the four-point Likert measure of the necessity of employment protection regulation, coded from 1 (“not necessary”) to 4 (“very necessary”). Column (1) restricts the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms. Column (2) restricts the sample to the *Partial Negation* and *Full Negation* arms. Cutpoints are the estimated thresholds between adjacent categories of the Likert scale. All specifications include country fixed effects and the full set of pre-treatment individual controls.

*** p<0.01, ** p<0.05, * p<0.1.

Table E.3: Wording effects on support for employment protection regulation (OLS estimates)

	<i>Support for employment protection regulation</i>			
	Direct vs. Partial Negation		Partial vs. Full Negation	
	Likert (1–4) (1)	Binary (0/1) (2)	Likert (1–4) (3)	Binary (0/1) (4)
Partial Negation	0.183*** (0.037)	0.077*** (0.019)		
Full Negation			0.022 (0.037)	0.015 (0.019)
AI Personal displacement risk	0.006 (0.038)	0.001 (0.020)	0.026 (0.038)	0.003 (0.019)
AI vulnerable group: Middle class	0.049 (0.041)	0.017 (0.021)	0.045 (0.040)	0.021 (0.020)
AI vulnerable group: Upper class	-0.179* (0.102)	-0.062 (0.049)	-0.057 (0.096)	-0.015 (0.048)
Female	0.018 (0.038)	0.022 (0.020)	0.014 (0.039)	0.012 (0.020)
Age	-0.011 (0.009)	-0.007 (0.005)	-0.006 (0.009)	-0.005 (0.005)
Complete HS education	0.091 (0.069)	0.050 (0.034)	0.112 (0.071)	0.051 (0.034)
More than HS education	0.057 (0.066)	0.038 (0.033)	0.133* (0.070)	0.065* (0.034)
Wealth index (1-10)	0.006 (0.012)	-0.000 (0.006)	-0.004 (0.012)	0.000 (0.006)
Currently employed	0.052 (0.042)	0.034 (0.022)	0.089** (0.043)	0.048** (0.022)
Political ideology (1-10)	-0.020** (0.008)	-0.010*** (0.004)	-0.016** (0.008)	-0.006* (0.004)
Collective action: neighbors act	0.048** (0.023)	0.025** (0.011)	0.040* (0.023)	0.016 (0.012)
State responsiveness	0.041 (0.026)	0.020 (0.013)	0.072*** (0.026)	0.022* (0.013)
Police response efficiency	0.036 (0.029)	0.017 (0.015)	0.009 (0.030)	0.002 (0.015)
Patience	0.006 (0.020)	0.002 (0.009)	0.007 (0.020)	0.010 (0.009)
Risk tolerance	-0.008 (0.019)	-0.003 (0.009)	-0.013 (0.020)	-0.008 (0.009)
Understands market competition	0.017 (0.039)	0.019 (0.020)	0.069* (0.040)	0.026 (0.020)
Understands price ceilings	-0.017 (0.038)	0.002 (0.019)	-0.013 (0.038)	0.008 (0.019)
Follows the news	0.018 (0.018)	0.005 (0.009)	0.022 (0.018)	0.013 (0.009)
Constant	2.510*** (0.232)	0.576*** (0.120)	2.446*** (0.233)	0.503*** (0.116)
Observations	2,634	2,634	2,634	2,634
R-squared	0.034	0.026	0.022	0.019

Notes: Each column reports OLS coefficients and robust standard errors in parentheses. The dependent variable in Columns (1) and (3) is the four-point Likert measure of the necessity of employment protection regulation, coded from 1 (“not necessary”) to 4 (“very necessary”). The dependent variable in Columns (2) and (4) is a binary indicator equal to one if the respondent reports that regulation is “somewhat” or “very” necessary. Columns (1) and (2) restrict the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms. Columns (3) and (4) restrict the sample to the *Partial Negation* and *Full Negation* arms. All specifications include country fixed effects and the full set of pre-treatment individual controls. *** p<0.01, ** p<0.05, * p<0.1.

Table E.4: Heterogeneous wording effects on support for employment protection regulation, by education

	<i>Support for employment protection regulation</i>	
	Direct vs. Partial Negation (1)	Partial vs. Full Negation (2)
<i>Education</i>		
Completed high school	0.047 (0.048)	0.064 (0.048)
More than high school	-0.001 (0.045)	0.097** (0.045)
<i>Treatments and interactions</i>		
Partial Negation	0.022 (0.058)	
Partial Negation $\times$ Completed HS	0.007 (0.068)	
Partial Negation $\times$ More than HS	0.087 (0.063)	
Full Negation		0.059 (0.059)
Full Negation $\times$ Completed HS		-0.027 (0.068)
Full Negation $\times$ More than HS		-0.061 (0.063)
Observations	2,634	2,634
R-squared	0.027	0.019

Notes: Each column reports OLS coefficients and robust standard errors in parentheses from linear probability models with interaction terms. The dependent variable is a binary indicator equal to one if the respondent reports that employment protection regulation is “somewhat” or “very” necessary. Column (1) restricts the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms; Column (2) restricts to respondents assigned to the *Partial Negation* and *Full Negation* arms. Main-effect coefficients on the treatment indicators (rows *Partial Negation* and *Full Negation*) capture the wording effect within the omitted education category (less than high school). Interaction terms capture the additional wording effect for each non-omitted education category, relative to the base. All specifications include country fixed effects and the full set of pre-treatment individual controls.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.5: Heterogeneous wording effects on support for employment protection regulation, by employment status

	<i>Support for employment protection regulation</i>	
	Direct vs. Partial Negation (1)	Partial vs. Full Negation (2)
<i>Employment</i>		
Employed	-0.007 (0.030)	0.068** (0.030)
<i>Treatments and interactions</i>		
Partial Negation	0.021 (0.033)	
Partial Negation $\times$ Employed	0.085** (0.040)	
Full Negation		0.040 (0.032)
Full Negation $\times$ Employed		-0.039 (0.040)
Observations	2,634	2,634
R-squared	0.027	0.019

Notes: Each column reports OLS coefficients and robust standard errors in parentheses from linear probability models with interaction terms. The dependent variable is a binary indicator equal to one if the respondent reports that employment protection regulation is “somewhat” or “very” necessary. Column (1) restricts the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms; Column (2) restricts to respondents assigned to the *Partial Negation* and *Full Negation* arms. Main-effect coefficients on the treatment indicators (rows *Partial Negation* and *Full Negation*) capture the wording effect within the omitted employment category (not employed). Interaction terms capture the additional wording effect for employed respondents, relative to the base. All specifications include country fixed effects and the full set of pre-treatment individual controls. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.6: Heterogeneous wording effects on support for employment protection regulation, by perceived AI vulnerability

	<i>Support for employment protection regulation</i>	
	Direct vs. Partial Negation (1)	Partial vs. Full Negation (2)
<i>AI vulnerable group</i>		
Middle class	-0.024 (0.029)	0.053* (0.028)
Upper class	-0.118* (0.070)	-0.007 (0.068)
<i>Treatments and interactions</i>		
Partial Negation	0.033 (0.028)	
Partial Negation $\times$ Middle class	0.083** (0.039)	
Partial Negation $\times$ Upper class	0.113 (0.097)	
Full Negation		0.045* (0.027)
Full Negation $\times$ Middle class		-0.063* (0.038)
Full Negation $\times$ Upper class		-0.014 (0.093)
Observations	2,634	2,634
R-squared	0.028	0.020

Notes: Each column reports OLS coefficients and robust standard errors in parentheses from linear probability models with interaction terms. The dependent variable is a binary indicator equal to one if the respondent reports that employment protection regulation is “somewhat” or “very” necessary. Column (1) restricts the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms; Column (2) restricts to respondents assigned to the *Partial Negation* and *Full Negation* arms. Main-effect coefficients on the treatment indicators (rows *Partial Negation* and *Full Negation*) capture the wording effect within the omitted vulnerability category (poor as most vulnerable). Interaction terms capture the additional wording effect for respondents who identify middle- or upper-class as most vulnerable, relative to the poor baseline. All specifications include country fixed effects and the full set of pre-treatment individual controls. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.7: Heterogeneous wording effects on support for employment protection regulation, by age

	<i>Support for employment protection regulation</i>	
	Direct vs. Partial Negation (1)	Partial vs. Full Negation (2)
<i>Age</i>		
30–44	-0.103*** (0.035)	-0.021 (0.034)
45–59	-0.146*** (0.037)	-0.044 (0.035)
60+	-0.130*** (0.050)	-0.082* (0.047)
<i>Treatments and interactions</i>		
Partial Negation	0.009 (0.035)	
Partial Negation $\times$ 30–44	0.091* (0.048)	
Partial Negation $\times$ 45–59	0.114** (0.050)	
Partial Negation $\times$ 60+	0.057 (0.067)	
Full Negation		0.026 (0.034)
Full Negation $\times$ 30–44		-0.048 (0.048)
Full Negation $\times$ 45–59		0.002 (0.048)
Full Negation $\times$ 60+		0.015 (0.067)
Observations	2,634	2,634
R-squared	0.027	0.020

Notes: Each column reports OLS coefficients and robust standard errors in parentheses from linear probability models with interaction terms. The dependent variable is a binary indicator equal to one if the respondent reports that employment protection regulation is “somewhat” or “very” necessary. Column (1) restricts the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms; Column (2) restricts to respondents assigned to the *Partial Negation* and *Full Negation* arms. Main-effect coefficients on the treatment indicators (rows *Partial Negation* and *Full Negation*) capture the wording effect within the omitted age category (18–29). Interaction terms capture the additional wording effect for each older age category, relative to the base. All specifications include country fixed effects and the full set of pre-treatment individual controls. *** p<0.01, ** p<0.05, * p<0.1.

Table E.8: Heterogeneous wording effects on support for employment protection regulation, by political ideology

	<i>Support for employment protection regulation</i>	
	Direct vs. Partial Negation (1)	Partial vs. Full Negation (2)
<i>Political ideology</i>		
Center-left	-0.070 (0.067)	0.125* (0.066)
Center	-0.157* (0.092)	0.036 (0.092)
Center-right	-0.232* (0.139)	0.010 (0.137)
Right	-0.263 (0.184)	0.012 (0.180)
<i>Treatments and interactions</i>		
Partial Negation	-0.030 (0.056)	
Partial Negation × Center-left	0.145** (0.071)	
Partial Negation × Center	0.113* (0.064)	
Partial Negation × Center-right	0.121 (0.073)	
Partial Negation × Right	0.105 (0.075)	
Full Negation		0.055 (0.055)
Full Negation × Center-left		-0.094 (0.069)
Full Negation × Center		-0.035 (0.062)
Full Negation × Center-right		-0.031 (0.072)
Full Negation × Right		-0.024 (0.073)
Observations	2,635	2,636
R-squared	0.025	0.019

Notes: Each column reports OLS coefficients and robust standard errors in parentheses from linear probability models with interaction terms. The dependent variable is a binary indicator equal to one if the respondent reports that employment protection regulation is “somewhat” or “very” necessary. Column (1) restricts the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms; Column (2) restricts to respondents assigned to the *Partial Negation* and *Full Negation* arms. Main-effect coefficients on the treatment indicators (rows *Partial Negation* and *Full Negation*) capture the wording effect within the omitted ideology category (Left). Interaction terms capture the additional wording effect for each non-omitted ideology category, relative to the base. All specifications include country fixed effects and the full set of pre-treatment individual controls. *** p<0.01, ** p<0.05, * p<0.1.

Table E.9: Heterogeneous wording effects on support for employment protection regulation, by gender

	<i>Support for employment protection regulation</i>	
	Direct vs. Partial Negation (1)	Partial vs. Full Negation (2)
<i>Gender</i>		
Female	0.045 (0.028)	0.006 (0.027)
<i>Treatments and interactions</i>		
Partial Negation	0.101*** (0.027)	
Partial Negation $\times$ Female	-0.050 (0.038)	
Full Negation		0.010 (0.026)
Full Negation $\times$ Female		0.009 (0.037)
Observations	2,634	2,634
R-squared	0.026	0.018

Notes: Each column reports OLS coefficients and robust standard errors in parentheses from linear probability models with interaction terms. The dependent variable is a binary indicator equal to one if the respondent reports that employment protection regulation is “somewhat” or “very” necessary. Column (1) restricts the sample to respondents assigned to the *Direct Wording* and *Partial Negation* arms; Column (2) restricts to respondents assigned to the *Partial Negation* and *Full Negation* arms. Main-effect coefficients on the treatment indicators (rows *Partial Negation* and *Full Negation*) capture the wording effect within the omitted gender category (male). Interaction terms capture the additional wording effect for female respondents, relative to the base. All specifications include country fixed effects and the full set of pre-treatment individual controls. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.