Discussion Paper Series

IZA DP No. 19062

November 2026

Artificial Intelligence and Support for Labor Market Policies: Evidence from a Survey Experiment

Fabio Angei
University of Cagliari and CRENoS

Silvia Balia
University of Cagliari,
CRENoS, HEDG and GLO

Rinaldo Brau
University of Cagliari and CRENoS

Giovanni Sulis
University of Cagliari,
CRENoS, LaborHub and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

IZA DP No. 19062 November 2026

Artificial Intelligence and Support for Labor Market Policies: Evidence from a Survey Experiment*

Abstract

We study whether alternative narratives about Artificial Intelligence (AI) shape workers' support for strengthening labor market policies and whether these responses differ by workplace AI use. We conduct a four-arm survey experiment on a representative sample of 2,105 Italian employees. Respondents are randomly assigned to a control group or to one of three information treatments framing AI as a technology that may substitute for human labor, complement workers' tasks, or increase the need for reskilling. We then elicit support for five labor market policy instruments, including both passive and active measures. Average treatment effects on policy support are small, but mask substantial heterogeneity by workplace AI use. Among non-users, the substitution narrative increases support for publicly funded training and hiring subsidies, while complementarity increases support for training. Among workplace AI users, both substitution and complementarity reduce support for unemployment insurance and short-time work schemes. Reskilling produces no statistically significant effects in either group. We interpret these heterogeneous effects using evidence that, among users, the substitution narrative increases the probability of expecting household financial losses from public intervention addressing AI's labor market effects, without a detectable change in perceived personal job-loss risk. Overall, workers' responses to the same information differ by workplace AI experience, both in the magnitude and direction of changes in support for specific labor market policies. The findings suggest that support for labor market protection may evolve as AI exposure and perceived vulnerability become more salient.

JEL classification

J65, J08, O33, D83, C93

Keywords

artificial intelligence, labor market policy, workers' perceptions, policy support, survey experiment

Corresponding author

Giovanni Sulis
gsulis@unica.it

* This research received funding from the Italian Ministry of University and Research (MUR) through the PRIN 2022 programme under the National Recovery and Resilience Plan (PNRR), project 2022YAPBL3 (CUP F53D23003210006), University of Cagliari. The study was registered in the AEA RCT Registry (Angei et al., 2025) and received ethical approval from the Ethical Commission of Università degli Studi di Cagliari on July 28, 2025 (approval number 0234435). We thank Marco Giovanni Nieddu, Milena Nikolova, Silvana Robone, and Mirco Tonin for comments and suggestions. We also thank participants at the 2026 SIEP Conference in Modena.

1 Introduction

Generative AI has diffused very rapidly among individuals and organizations (Stanford Institute for Human-Centered Artificial Intelligence, 2026). Its implications for workers and firms depend on how it is deployed: AI may automate existing tasks, potentially reducing labor demand, but it may also support workers' decision-making and raise their productivity (Autor, 2024; Acemoglu et al., 2026). Field evidence from customer-support workers illustrates its potential to complement labor: Brynjolfsson et al. (2025) find that generative AI raises productivity, with particularly large gains among less experienced and lower-skilled workers. Changes in workplace practices and productivity need not, however, translate quickly into changes in earnings or hours worked. Using Danish data, Humlum and Vestergaard (2026) document rapid workplace adoption, perceived productivity gains, and substantial task reorganization following the introduction of ChatGPT, while effects on workers' earnings and hours remained limited over the first two years.

Changes in tasks and workplace practices raise the question of how labor market policies can protect workers against potential displacement while helping them to develop the skills needed to work with AI and benefit from it. Income protection and support for workers facing job loss may be relevant where AI reduces labor demand, while training and other active labor policies can help workers respond to changing tasks and skill requirements. Early evidence using vacancy data shows that AI adoption is associated with changing skill requirements within firms (Acemoglu et al., 2022). More recent evidence documents a growing demand for new skills, particularly in IT and AI-related fields, as well as substantial cross-country differences in emerging skill gaps (Jaumotte et al., 2026). Consistent with this evidence, the OECD emphasises upskilling, reskilling, and lifelong learning as central margins of adjustment to AI-driven technological change (OECD, 2026).

In this paper, we study how information about the implications of AI diffusion in the workplace affects workers' views and support for specific labor market policies ranging from passive insurance to active interventions. Understanding workers' demand for labor market policies is an important complement to identifying their optimal design: policies that lack public support may be difficult to adopt or sustain. A growing research stream measures workers' and occupations' technical exposure to AI (Felten et al., 2021; Eloundou et al., 2024), while related works examine the channels through which AI may automate, complement, or reshape human work (Acemoglu and Restrepo, 2018; Autor,

2024). Yet, technical exposure does not mechanically translate into support for a particular policy response. Recent research has begun to examine how automation and AI shape policy preferences, but we still know relatively little about how workers respond to competing narratives about AI across a common menu of labor market instruments, and whether these responses differ with prior workplace experience with AI.[1] Workers who interpret AI primarily as a displacement risk may favor income protection or employment-preserving instruments, whereas those who view AI as complementary to labor or as challenging skill requirements may place greater weight on training and adaptation. Familiarity with, and use of AI for work may shape workers' perceptions of its potential effects and their responses to information about those effects. Consistent with this possibility, U.S. survey evidence shows that AI users are more likely than non-users to expect workplace AI to affect their future job opportunities (Pew Research Center, 2025). Studying how workers translate alternative AI narratives into policy preferences, and how these responses vary with workplace AI use, therefore complements the technical analysis of AI exposure.

We conduct a survey experiment on a representative sample of Italian employees aged 18–60. All respondents receive a common definition of AI and are then randomly assigned to one of four experimental groups. The three treatment groups receive an additional short narrative alternatively describing AI as a technology that may substitute for human labor, complement workers' tasks, or increase the need for reskilling and continuous training, while the control group receives no further information. The treatment narratives were designed to reflect distinct, potentially coexisting implications of AI for workers, drawing on research on the Italian labor market. Specifically, Dalla Zuanna et al. (2024) document heterogeneous occupational exposure to AI and distinguish between its substitution and complementarity dimensions, informing the corresponding narratives. Ferri et al. (2024) emphasise the role of continuous training and professional updating in enabling workers to adopt and integrate AI into their daily tasks, informing the reskilling narrative.[2] These studies provide an empirical grounding for the treatments and anchor their content in the national

[1] Recent work by Chen et al. (2025) finds that AI may have heterogeneous effects on labor demand and skill requirements across jobs and occupations suggesting both a substitutability and a complementarity relation. The latter emerges in particular in tasks and jobs that require high human-AI collaboration.

[2] Evidence from other countries is consistent with this role of training. Lane et al. (2023) find that workers who receive training to work with AI are significantly more likely to report improvements in job performance and working conditions. Humlum and Vestergaard (2026) document higher workplace AI chatbot adoption among workers receiving employer-provided training, with adoption highest when training is combined with employer encouragement and access to enterprise chatbots.

context of our experiment.

Following the experimental assignment, we measure support for five labor market policy instruments: unemployment benefits, short-time work schemes, minimum-income policies, publicly funded worker training, and hiring subsidies for firms. These instruments differ not only in their objectives but also in their financing. We also examine perceived job-loss risk, the expected financial impact of AI-related public intervention on the respondent's household, the perceived usefulness of public intervention, and willingness to support a petition for institutional action on AI regulation.

Our preliminary descriptive analysis shows that support for labor-market policies is systematically associated with respondents' geographic location and public-sector employment, while self-assessed AI knowledge is positively associated with support across all five policy instruments. The survey experiment provides little evidence of sizable average effects of the three AI narratives on support for strengthening different active and passive labor-market policies. These limited average effects, however, mask substantial heterogeneity by actual workplace AI use. Among non-users, the substitution narrative increases support for publicly funded training and hiring subsidies, while complementarity increases support for training. Among AI users, both narratives reduce support for unemployment insurance and short-time work schemes. The reskilling narrative produces no statistically significant effects in either group. These patterns suggest that workplace AI experience shapes how workers translate information about its labor market implications into policy preferences. Additional evidence on perceptions and real-stake behavior further underscores this role. Among AI users, substitution increases by 6.2 percentage points the probability of expecting public intervention to leave their household financially worse off, while its effect on perceived personal job-loss risk is close to zero: greater displacement is associated with higher spending on passive labor market protection without a corresponding increase in the perceived likelihood of needing such protection; this makes these programs less attractive. Among non-users, both substitution and complementarity increase petition support (7.5 percentage points). Together with their positive effects on support for publicly funded training, this pattern is consistent with AI information making adaptation and institutional action more salient among workers without direct workplace experience with the technology.

This paper is related to three strands of literature. First, it relates to survey-experimental research on the formation of policy preferences. Support for policy responses to issues of broad

public concern depends on perceived effectiveness, incidence, fairness, self-interest, and beliefs about the appropriate role of government (Stantcheva, 2021; Haaland et al., 2023; Dechezleprêtre et al., 2025). A recurring finding in this literature is that information can alter beliefs about the economic condition without producing large average changes in support for redistributive policies (Kuziemko et al., 2015; Engelhardt and Wagener, 2018). At the same time, information may shift support across specific policy instruments even when general demand for government intervention in the economy changes little (Lergetporer et al., 2018; Di Tella and Rodrik, 2020). This distinction is particularly relevant in the labor market, where policy instruments differ in their incidence, target groups, financing, and balance between insurance, protection, and active policy adjustment.[3]

Second, our paper speaks to the growing literature on how automation and AI shape workers' economic concerns and policy preferences. Earlier work links exposure or perceived vulnerability to automation to preferences for redistribution and social insurance (Thewissen and Rueda, 2019), and distinguishes between compensatory policies and policies aimed at protecting workers from technological change (Gallego et al., 2022). Related evidence shows that workers differentiate among alternative responses to technological change, including compensation, social investment, and restrictions on automation (Busemeyer et al., 2023; Jeffrey, 2021). Using a discrete-choice experiment, Cattaneo et al. (2025) show that workers are willing to accept lower wages in exchange for lower automation risk, with substantial heterogeneity across demographic groups. Related experimental evidence shows that information about occupation-specific automation risk can shift the type of response workers favor: Golin and Rauh (2022) find limited effects on intentions to retrain or switch occupations, but greater willingness to seek employment protection through union membership. A recent strand of experimental work focuses specifically on generative AI. Haslberger et al. (2026) experimentally expose respondents to generative AI in realistic work tasks and show that such exposure does not increase perceived labor market risk but strengthens support for activating social policies. Most closely related to our setting, Shen et al. (2026) provide Chinese white-collar workers with information about occupational exposure to generative AI. Their displacement-risk treatment increases perceived job vulnerability and support for unemployment insurance, universal

[3]Two papers are especially related to our contribution. Cotofan and Matakos (2025) analyze the relationship between labor market shocks and preferences for alternative labor market and redistributive policies. Duval et al. (2024) study how beliefs and individual endowments shape support for employment protection legislation and whether new information changes such support.

basic income, and publicly funded retraining. Our studies differ along two important dimensions. First, we use a nationally representative sample of Italian employees, rather than a local sample restricted to white-collar workers. More importantly, our experimental design does not simply vary the salience of displacement risk. We compare three qualitatively different interpretations of the role of AI - substitution, complementarity, and reskilling - and examine how they impact support for strengthening a specific set of passive and active labor market policies. This allows us to study not only whether information changes the level of policy support, but also the composition of labor market policy demand.[4]

Third, we connect to evidence showing that workers' responses to technological change depend on their own exposure to automation and new technologies. Policy preferences differ systematically with workplace technology use (Busemeyer et al., 2023), while broader evidence shows that exposure to new technologies can generate heterogeneous labor market outcomes across workers (Ge and Zhou, 2020) and can affect political attitudes and behavior (Anelli et al., 2021). This perspective is especially relevant for AI. Adoption remains highly uneven across workers, firms, and countries (Humlum and Vestergaard, 2026; Bick et al., 2026a), and recent evidence documents substantial differences in realized GenAI adoption even among workers performing similar tasks (Bick et al., 2026b). We build on this literature by examining whether actual workplace AI use changes how workers translate alternative randomized narratives into support for specific labor market policies.

Our study is also relevant to the broader literature on the determinants of labor market institutions and regulation. Aghion et al. (2011) emphasize the role of beliefs and trust, arguing that the quality of labor relations may influence the demand for labor market regulation. Algan and Cahuc (2009) show that, in countries where civic virtues are stronger, individuals tend to prefer unemployment benefits to employment protection legislation as a means of insuring workers against job loss. Alesina et al. (2015) highlight the role of family values in shaping preferences for labor regulation, while Botero et al. (2004) emphasize the importance of countries' deep institutional roots.[5] More recently, Acemoglu et al. (2026) show that automation can reduce labor's share of national income and thereby increase the threat of social conflict. A capitalist state may respond

[4]Related work shows that alternative positive and negative frames about AI can also affect beliefs and policy support in other domains (see, e.g., Brodeur et al., 2025).

[5]For the Italian context, Cardullo et al. (2024) show that current local differences in cooperative labor relations are partly rooted in historically more positive attitudes towards non-family members, greater participation in community life, and a stronger willingness to join trade unions.

by limiting automation or, if automation continues, by addressing the resulting political tensions through redistribution toward workers or repression of political action. We contribute to this literature by shifting attention from the strategies available to the state and capital owners to the preferences of workers themselves. In particular, we provide micro-level evidence on how workers evaluate alternative labor market policy instruments when confronted with rapid technological change.

The remainder of the paper proceeds as follows. Section 2 describes the survey design, randomized treatments, and outcomes. Section 3 presents the empirical strategy. Section 4 reports the empirical results, beginning with a descriptive benchmark for policy support and then examining average treatment effects, heterogeneity by workplace AI use, additional outcomes, and robustness checks. Section 5 discusses the interpretation of the findings, possible mechanisms, and the potential evolution of policy demand as AI diffuses further. Section 6 concludes. Additional material is gathered in the Appendix.

2 Survey Design

2.1 Sample and data quality

Our analysis relies on an original survey of Italian employees aged 18–60, that took place in October 2025 using CAWI (Computer-Assisted Web Interviewing). Respondents were recruited from *IPSOS* online panel using quota sampling to match the composition of the Italian employee population along four dimensions: gender, age, education, and macro-region.[6] Quotas distinguished between women and men; five age groups (18–24, 25–34, 35–44, 45–54, and 55–60); respondents with and without a university degree; and four macro-regions (North West, North East, Center, and South with Islands). The questionnaire collected demographic and socio-economic characteristics, information on respondents' current employment situation, trust in the Italian state, self-assessed AI knowledge, workplace AI use, and perceived job-loss risk.During fieldwork, *IPSOS* independently administered quality-control procedures to identify potentially inattentive or inconsistent respondents. These

[6] *IPSOS* is a leading company specialized in survey research. It compensates members of its respondent panel for completing the survey through monetary payments, vouchers, or charitable donations. Incentives are of equivalent monetary value across respondents, regardless of the form of compensation. These incentives are linked to survey participation and not to the content or accuracy of responses to individual questionnaire items.

included checks for implausible response patterns, such as inconsistencies with reported employment status or the selection of multiple unlikely events, and survey completion times. Respondents completing the questionnaire in less than one-third of the median duration were automatically excluded and replaced. Out of 2,123 completed interviews, we retained respondents with complete baseline information, yielding a full sample of 2,105 observations.

Table 1, column (1), reports the main descriptive statistics for the full sample. Respondents are 44 years old on average; 46% are women, 28% hold at least a university degree, and 53% live in Northern Italy. Most respondents have a permanent contract, and on average they work about 35 hours per week. Familiarity with AI and workplace use are widespread: 49% of respondents report a high level of knowledge of AI and its potential applications in everyday life. On the other hand, approximately 37% report using AI for work frequently, either daily or a few times a week. This share rises to 61% when respondents who use AI rarely are also included, as reported in the Table.[7]

Figure 1 compares the demographic and geographic composition of our sample with the 2024 Italian Labor Force Survey. The figure displays percentage-point deviations of sample shares from the corresponding LFS benchmarks. Our sample closely matches the official national survey data in terms of demographic and geographic characteristics, while slightly over-representing highly educated workers and middle-aged workers.

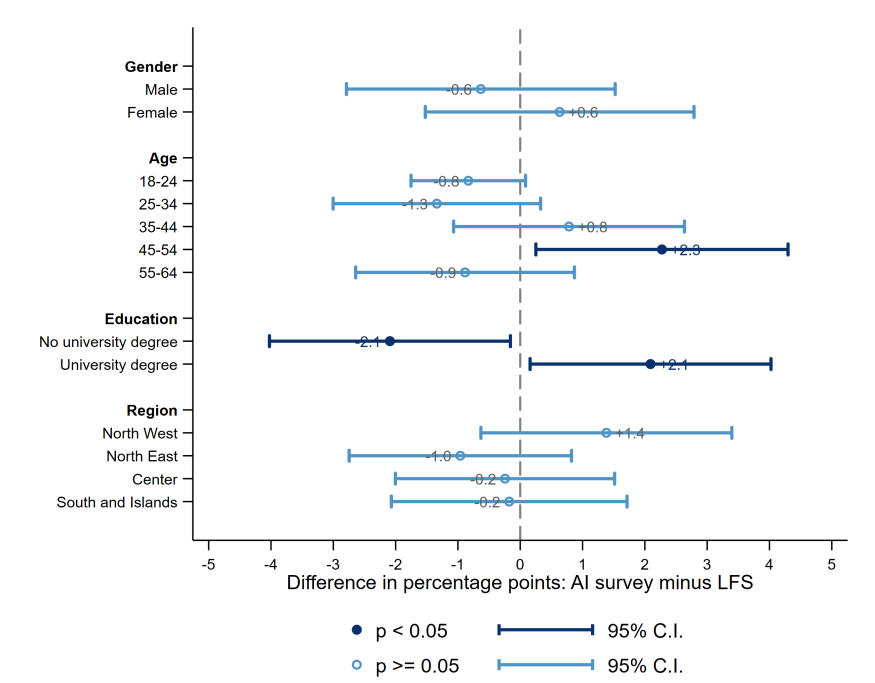
2.2 Randomized AI information treatments

Our experiment examines how information emphasizing AI's potential to substitute for human labor, complement workers' tasks, or increase the need for reskilling shapes workers' labor market policy preferences.[8] We randomly assign respondents to one of three information treatments or a control group, as shown in Figure 2, which summarizes the structure of the survey. Before assignment, all respondents read a common definition of AI as technologies and computer programs that analyze large amounts of data and perform tasks typically requiring human intelligence, including

[7]For comparison, recent surveys suggest a large workplace use of AI in the Italian public sector. In a survey conducted by the Italian Department of Public Administration between November 2025 and January 2026, 43.1% of respondents reported using generative AI tools almost every day and a further 32.1% about once a week. Similarly, FPA's 2026 PanelPA survey of 500 public employees found that 66% used AI tools every day or every week, while 23% used them rarely and 11% never did. See the Department of Public Administration survey and the FPA PanelPA 2026 survey. Although these surveys focus only on the public sector and use a different sampling design relative to our survey, they similarly document a large workplace AI diffusion.

[8]In Appendix A we describe in detail the policies and their financing in the Italian institutional context.

Figure 1: Survey sample compared with Italian employees in the 2024 Labor Force Survey



Notes: The figure reports percentage-point differences between the survey sample ($N = 2{,}105$) and employees in the 2024 Italian Labor Force Survey ($N = 128{,}740$). The LFS estimates are weighted using the survey population weights, while our survey estimates are unweighted. Positive values indicate that a group is overrepresented in our survey relative to the LFS. Horizontal bars report 95% confidence intervals for the differences. Dark markers denote differences statistically significant at the 5% level. The LFS sample covers employees aged 15–64, while the survey covers respondents aged 18–60. Comparisons for the youngest and oldest age groups are therefore approximate because the LFS provides age in grouped categories.

Table 1: Sample characteristics

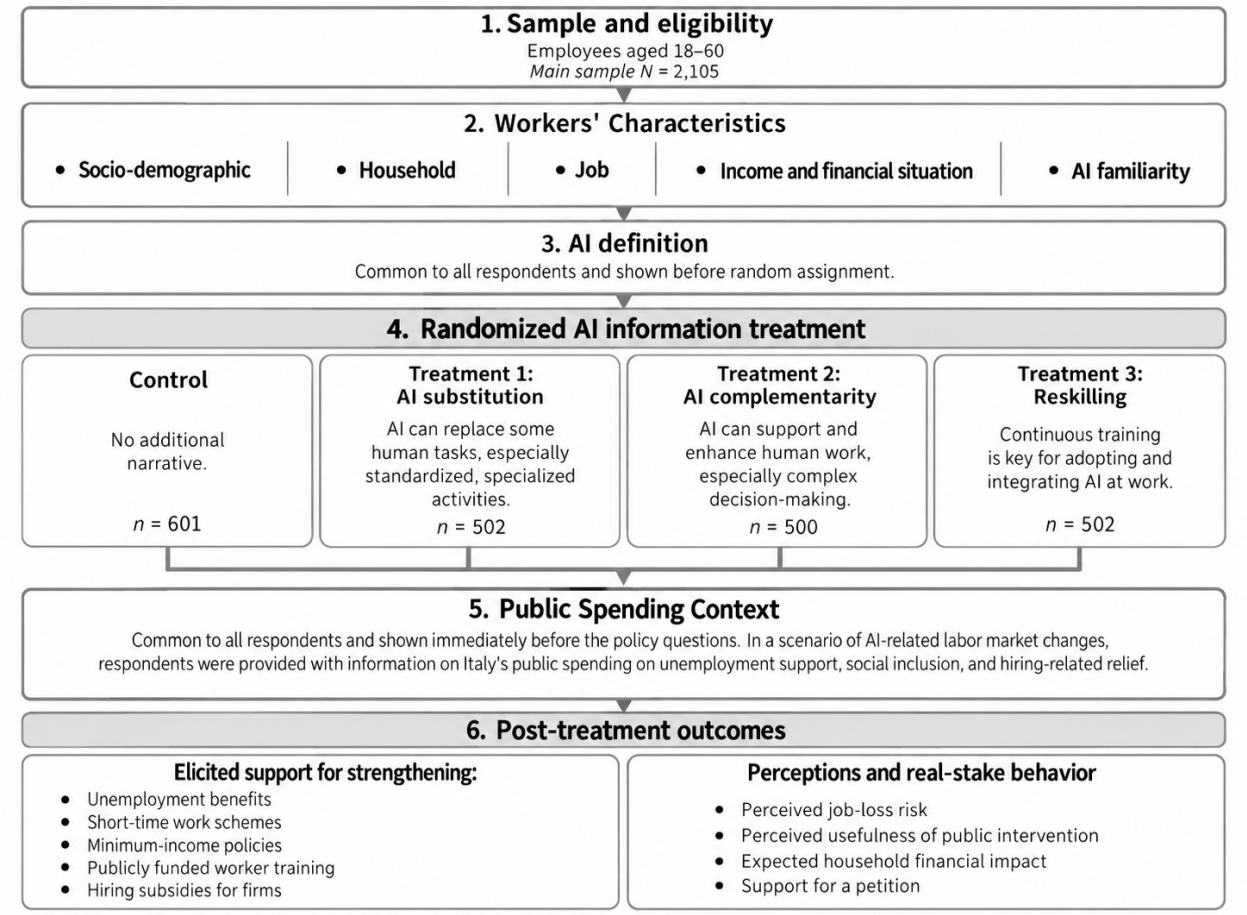
	Full sample (1)	Control (2)	AI substitution (3)	AI complementarity (4)	Reskilling (5)
Panel A. Socio-demographic characteristics					
Age	44.160	44.715	43.755	44.466	43.594
Female	0.459	0.479	0.412	0.500	0.440
University degree or higher	0.276	0.271	0.267	0.292	0.275
North	0.533	0.524	0.542	0.522	0.544
Center	0.210	0.225	0.191	0.222	0.197
South and Islands	0.258	0.251	0.267	0.256	0.259
Panel B. Household characteristics					
Household size	2.865	2.905	2.786	2.938	2.826
Lives with partner	0.688	0.696	0.655	0.728	0.671
Has children	0.579	0.622	0.548	0.586	0.550
Number of children	0.914	1.005	0.847	0.934	0.855
Panel C. Job characteristics					
Public sector	0.203	0.209	0.217	0.203	0.183
Part-time	0.200	0.196	0.193	0.230	0.181
Permanent contract	0.840	0.872	0.817	0.838	0.827
Temporary contract	0.120	0.085	0.143	0.122	0.137
Workplace has more than 15 employees	0.693	0.701	0.723	0.678	0.666
Days worked per week	5.156	5.131	5.175	5.172	5.149
Hours worked per week	35.314	35.258	35.072	34.642	36.300
Current job tenure (years)	11.630	12.397	11.017	11.287	11.654
Panel D. Income and financial situation					
Low personal income	0.144	0.136	0.139	0.177	0.124
Medium-low personal income	0.265	0.247	0.254	0.263	0.301
Medium-high personal income	0.384	0.413	0.419	0.334	0.365
High personal income	0.207	0.205	0.188	0.225	0.209
Make ends meet	3.274	3.253	3.241	3.268	3.338
Panel E. AI familiarity and pre-treatment belief					
Uses AI at work	0.610	0.622	0.614	0.612	0.591
High AI knowledge	0.489	0.485	0.482	0.496	0.494
Perceived risk of job loss before treatment	2.395	2.393	2.351	2.396	2.440
Observations	2,105	601	502	500	502

Notes: The table reports sample means for the full sample and each experimental arm. For binary variables, means represent sample shares. Workplace size equals one for workplaces with more than 15 employees and is missing for respondents who do not report their workplace size. Personal income refers to net monthly income, grouped into four categories: low, medium-low, medium-high, and high. Make ends meet is measured on a six-point scale, ranging from 1 ("with great difficulty") to 6 ("very easily"). Uses AI for work equals one for respondents who report using AI for work every day, several times per week, or rarely, and zero for those who report never using it. High AI knowledge equals one for respondents who report knowing AI "fairly well"or "very well", and zero for those who report knowing it "not at all"or "a little". Both indicators are missing for respondents who prefer not to answer. Pre-treatment perceived risk of job loss is measured on a five-point agreement scale about whether respondents might lose their job within the next 36 months, ranging from 1 ("strongly disagree") to 5 ("strongly agree").

understanding and generating text, recognizing images or sounds, making predictions and decisions, and proposing solutions to problems. The control group receives only this definition, while each treatment group reads an additional narrative.

All three narratives draw on recent studies on the Italian labor market, and explicitly refer to a

Figure 2: Survey flow and randomized AI information treatments



Notes: The figure summarizes the survey structure for the 2,105 respondents in the final sample. Respondents first complete the eligibility screen and baseline module, then are randomly assigned to the control group ($N = 601$), the AI-substitution treatment ($N = 502$), the AI-complementarity treatment ($N = 500$), or the reskilling treatment ($N = 502$). The post-treatment module measures support for five labor market policies and the additional outcomes examined in the paper.

"recent scientific analysis" to signal that the information was grounded in scientific evidence. The substitution and complementarity narratives draw on Dalla Zuanna et al. (2024), while the reskilling frame builds on evidence from Ferri et al. (2024). The narratives use a neutral tone and closely parallel wording to ensure comparability across treatments. In particular, the substitution and complementarity frames share identical introductory wording and differ only in the final statement about AI's implications for human labor. The substitution treatment informs the respondent that AI can replace some human tasks, especially standardized, specialized activities. The complementarity treatment emphasizes how AI can support and enhance human work, especially in complex decision-making. The reskilling treatment highlights the importance of continuous training for adopting and

integrating AI into daily work.[9]

Random assignment resulted in 601 respondents in the control group, 502 in the substitution group, 500 in the complementarity group, and 502 in the reskilling group within the full sample. Table E1 in the Appendix shows no systematic imbalance in treatment assignment across socio-economic characteristics or actual AI use for work.[10] Importantly for the heterogeneity analysis, pre-treatment workplace AI use is balanced across treatment arms.

2.3 Support for strengthening labor market policies

Our information experiment examines whether salient qualitative information about AI's labor market implications affects workers' support for strengthening the five labor policy instruments. The latter are unemployment insurance, short-work programs, minimum income schemes, publicly funded training and hiring subsidies for firms and are described in more detail in Appendix A. Focusing on specific policy instruments, that could plausibly be used to address the effects of AI on work, rather than public intervention in general allows us to examine whether the treatments change the composition of labor market policy demand.

Regardless of treatment assignment, all respondents receive the same information public expenditures on unemployment support, social inclusion, and tax and social-contribution relief for hiring in Italy, based on the most recent data available at the time of the survey. They are then asked to indicate their support for strengthening five policy instruments in light of recent and anticipated changes in the labor market associated with AI adoption.[11] Respondents rate their support for strengthening each policy on a five-point scale ranging from 1 ("not at all in favor") to 5 ("strongly in favor"). The five policies are presented in a random order across respondents.[12]

Hiring subsidies, publicly funded training, and unemployment insurance receive the highest

[9] Appendix D reports close English translations of the original Italian treatment texts.

[10] Three covariate–arm associations are statistically significant at the 5% level: male gender is positively associated with assignment to the substitution arm (4.7 percentage points), living with a partner with assignment to the complementarity arm (5.6 percentage points), and having a temporary contract is negatively associated with assignment to the control arm (6.6 percentage points). Given the number of covariate–arm comparisons, these isolated imbalances are consistent with chance variation under random assignment.

[11] The fiscal context is described as follows: "In 2023, Italy devoted approximately 3.4% of GDP, or EUR 73 billion, to unemployment support, social inclusion, and tax and social-contribution relief for hiring." The question reads: "In light of recent and upcoming changes in the labor market linked to the use of Artificial Intelligence, how favorable would you be to strengthening each of the following public interventions?".

[12] Table D1, Panel A, in the Appendix reports definitions and descriptive statistics for the five policy-support outcomes.

Table 2: Preferences for strengthening labor market policies

Group	Category	N	Unemployment insurance	Short-time work	Minimum income	Publicly funded training	Hiring subsidies
All sample							
	Control	601	62.1	54.1	47.1	67.9	72.7
	AI substitution	502	54.4	54.0	46.0	67.3	72.1
	AI complementarity	500	53.8	50.6	46.4	66.4	69.4
	Reskilling	502	58.6	55.4	48.0	69.3	72.7
Control group							
Gender	Male	313	62.3	56.9	48.6	70.3	72.2
	Female	288	61.8	51.0	45.5	65.3	73.3
Age	18-34	115	60.9	48.7	48.7	60.9	69.6
	35-49	260	60.0	52.7	47.3	65.8	69.6
	50-60	226	65.0	58.4	46.0	73.9	77.9
Education	No university degree	438	65.1	54.8	45.2	69.4	74.2
	University degree	163	54.0	52.1	52.1	63.8	68.7
Region	North	315	61.6	55.2	43.2	67.6	70.2
	Center	135	61.5	50.4	45.9	68.1	70.4
	South and Islands	151	63.6	55.0	56.3	68.2	80.1

Notes: The table reports shares for 0–1 indicators by experimental arm and, within the control group, by sociodemographic group; reported values are multiplied by 100. Each policy support outcome equals one for responses 4 or 5 on the original five-point scale. Experimental-arm rows use the full main regression sample; control-group rows describe the 601 respondents assigned to the control group.

levels of support. As shown in the top panel of Table 2, approximately 73% of respondents in the control group express high support for strengthening hiring subsidies, followed by almost 68% for publicly funded training and 62.1% for unemployment insurance. Support is lower for short-time work schemes (54.1%) and the minimum income scheme (47.1%). The ordering of policies by the share of respondents expressing high support is unchanged across experimental arms. The largest differences relative to the control group concern unemployment insurance: approximately 54% of respondents in the substitution and complementarity groups express high support, compared with 62% in the control group. Differences across arms are smaller for the other policy instruments.

Table 2 also describes variation across demographic groups within the control group (see bottom panel). Workers aged 50–60 are more supportive of training, hiring subsidies, and short-time work schemes than younger respondents. University graduates express lower support for unemployment insurance but higher support for the minimum income for respondents without a degree. Geographical differences are most pronounced for the minimum income and hiring subsidies, both of which receive the highest support in the South and Islands. Support for the remaining policies varies less across macro-regions.

2.4 Workplace AI use

In the empirical analysis, we exploit respondents' self-reported workplace experience with AI, which we measure in the survey before the treatment.[13] Experience with AI for work may influence how workers interpret information about AI's labor market implications and, in turn, their support for specific public policies.[14] This motivates our analysis of treatment-effect heterogeneity by workplace AI use. We first compare users and non-users and examine treatment-arm balance within each group. Table B1 in the Appendix documents substantial differences between workplace AI users and non-users. Users are younger on average (42.6 versus 46.6 years), more likely to hold a university degree (35.7% versus 15.0%), and more likely to report high AI general knowledge (64.4% versus 24.7%). They are also more likely to work in the public sector, less likely to work part-time, and more likely to be in the highest personal-income category. By contrast, the difference in perceived job-loss risk between AI users and non-users as measured before treatment is not statistically significant.

Table B2 and Table B3 in the Appendix report respondent characteristics respectively for AI users and non-users. Treatment assignment was randomized across the full sample and not stratified by workplace AI use. Randomization therefore ensures balance in expectation within these pre-treatment subgroups, but additional imbalances may still arise. We detect no statistically significant differences across arms for most characteristics. At the 5% level, exceptions concern gender and income composition among non-users, and household characteristics and income composition among users.

3 Empirical Strategy

We estimate the causal effect of being exposed to qualitative information about the implications of AI in the labor market on workers' support for strengthening each of the five policy instruments described in Section 2. Specifically, we estimate the following regression model using Ordinary Least Squares (OLS):

$$Y_{ij} = \alpha_j + \beta_{1j}\,\text{Substitution}_i + \beta_{2j}\,\text{Complementarity}_i + \beta_{3j}\,\text{Reskilling}_i + \varepsilon_{ij}, \tag{1}$$

[13] We classify as workplace AI users respondents who report using AI for work every day, several times a week, or rarely, and as non-users those who report never using AI for work.

[14] Most importantly, AI use for work is different from the self-assessed general AI knowledge that we discuss in the following parts of the paper.

where Y_{ij} is an ordinal variable measured on a five-point scale indicating respondent i's support for strengthening policy instrument j. The coefficients β_{1j}, β_{2j}, and β_{3j} capture the average causal effects of being assigned and shown the substitution, complementarity, and reskilling narratives, respectively, relative to the control group. This specification follows directly from the randomized design and does not condition on post-treatment variables.

Moreover, we explore potential heterogeneity of the treatment effects by workplace AI use, as workers who use AI for work may respond differently to the same narrative than those who do not. We estimate an extended version of equation (1), which includes a binary indicator for workplace AI users alongside its interactions with the treatment indicators. This allows users and non-users to differ both in their mean support in the control group and in their responses to the information treatments. The regression model is the following:

$$Y_{ij} = \alpha_j + \sum_{k \in \{S,C,R\}} \beta_{kj} T_{ki} + \gamma_j \,\text{AI User}_i + \sum_{k \in \{S,C,R\}} \delta_{kj} \left(T_{ki} \times \text{AI User}_i\right) + \varepsilon_{ij}, \tag{2}$$

where T_{ki} indicates assignment to treatment narrative k (substitution, complementarity, or reskilling). In this specification, β_{kj} captures the average effect of each treatment k among non-users, while $\beta_{kj} + \delta_{kj}$ captures the corresponding effect among AI users.

4 Results

We first document how support for strengthening the five labor market policies varies with respondents' characteristics in the full sample. We then estimate the average effects of the three AI narratives and examine heterogeneity by workplace AI use. To help interpret these responses, we also examine treatment effects on perceived job-loss risk, the expected financial impact of public intervention on respondents' households, its perceived usefulness, and willingness to support institutional action on AI. These additional outcomes inform the discussion in Section 5. The section concludes with a summary of robustness checks, reported in full in the Appendix.

4.1 What drives policy support?

Figure 3 reports average marginal effects from separate probit models for each binary policy support outcome. The models include a rich set of socio-demographic, household, workplace, and income characteristics, together with trust in the state, pre-treatment perceived risk of job loss, AI knowledge, and workplace AI use. All specifications include treatment-arm indicators.

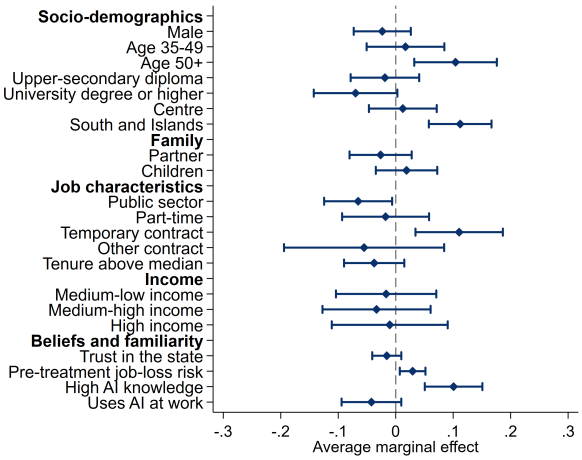
Age and geographic area of residence are systematically related to policy support. Compared with workers aged 18–34, those aged 50–60 are more likely to express high support for strengthening all policy instruments except the minimum income scheme. Workers aged 35–49 also show greater support for training relative to the youngest group, although this difference is significant only at the 10% level. Respondents living in the South and Islands are more supportive than those in the North of all five interventions, with the difference for short-time work significant only at the 10% level. Men are less likely than women to support hiring subsidies. University graduates show lower support for unemployment insurance than respondents with lower-secondary education or less, although only at the 10% level. There is also little evidence of differences across income categories relative to the lowest income group, or by partnership status and the presence of children in the household.

When considering job characteristics, workers in a temporary contract are more likely than those in a permanent contract to support unemployment insurance. Public-sector workers show lower support for unemployment insurance and hiring subsidies than private-sector workers, and lower support for training at the statistical 10% level. Workers with firm tenure above the sample median are less likely than those at or below the median to support the minimum income and hiring subsidies, both differences are significant only at the 10% level.

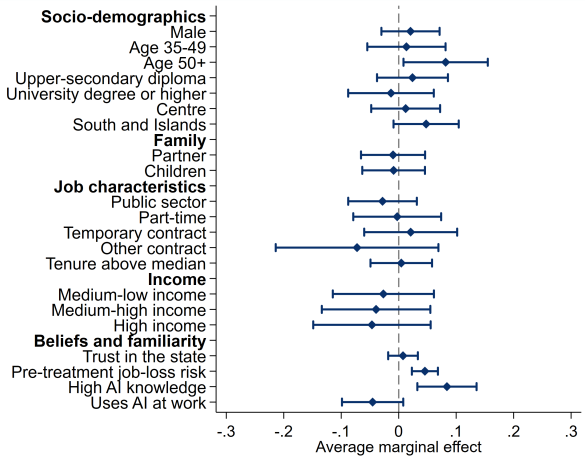
High self-assessed AI knowledge predicts greater support for all five policies, with average differences of approximately 7–11 percentage points relative to respondents without high self-assessed knowledge. Workplace AI users, however, are less likely than non-users to support short-time work and training, although these differences are significant only at the 10% level. Greater pre-treatment perceived risk of job loss predicts higher support for unemployment insurance, short-time work, and the minimum income schemes, with weaker evidence for hiring subsidies at the 10% level. Higher trust in the state predicts lower support for the minimum income and higher support for hiring subsidies.

Figure 3: Analysis of the support for the labor market policies

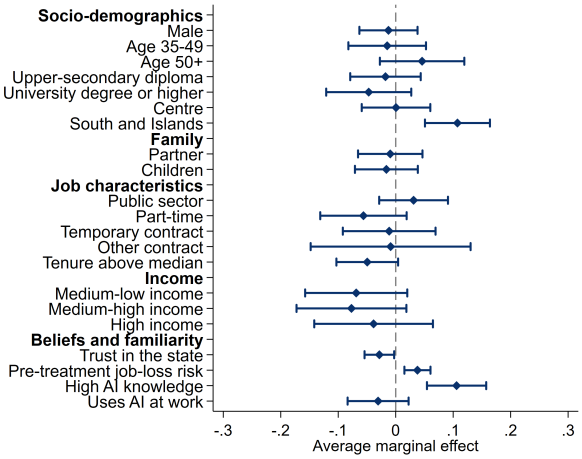
(a) Unemployment insurance



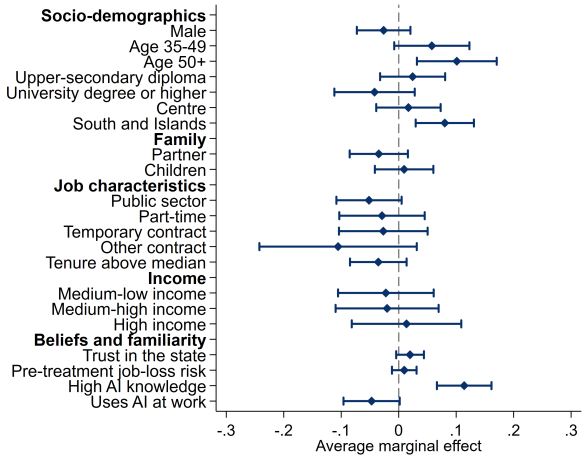
(b) Short-time work



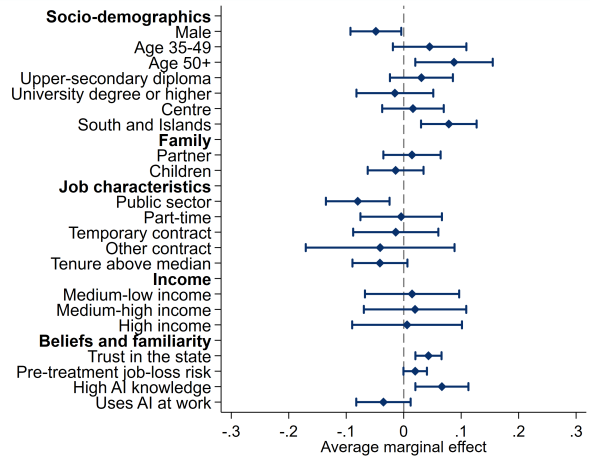
(c) Minimum income



(d) Publicly funded training



(e) Hiring subsidies



Notes: The figure reports average marginal effects from separate probit models for each policy-support outcome. Each outcome equals one for responses 4 or 5 on the original five-point scale, ranging from 1 ("not at all in favor") to 5 ("strongly in favor"), and zero otherwise. All models include treatment-arm indicators and use the same sample of 1,759 respondents with complete information on all included covariates. Horizontal bars report 95% confidence intervals based on heteroskedasticity-robust standard errors.

4.2 Perceptions and real-stakes behavior

To better characterize workers' perceptions of labor-market changes and public intervention addressing the effects of AI, we examine three outcomes measured after the experimental treatments: (a) perceived usefulness of public intervention in governing AI's impact on the labor market, capturing beliefs about its social value; (b) expected household financial impact, based on whether respondents expect public intervention addressing AI's labor-market effects and the associated public expenditure to leave their household financially better or worse off; and (c) perceived individual risk of job loss, capturing respondents' assessments of the possibility of losing their job.[15] We complement these measures with a real-stakes behavioral outcome: respondents' willingness to support a petition calling for institutional action to address the effects of AI on firms and workers.[16] Respondents are informed that their answers remain anonymous and that the aggregate share supporting the petition will be communicated to the Observatory on AI Adoption in the Labor Market at the Italian Ministry of Labor and Social Policies. We therefore interpret petition support as an action-oriented outcome capturing respondents' willingness to endorse a collective call for institutional intervention.[17] Communicating aggregate support to a public institution gives respondents' choices a consequence beyond the survey setting.

We first describe how these outcomes vary across treatment arms and respondent characteristics. Perceived usefulness of public intervention and perceived job-loss risk vary little across treatment arms, while expected household financial loss is more common in the substitution arm (top panel of Table 3). More precisely, in each experimental group, roughly two-thirds of respondents regard public intervention as useful, while 11–15% expect it to leave their household financially worse off and 13–16% perceive a risk of job loss. Petition support is widespread across experimental groups, ranging from 78.4% in the control group to 80.6% in the AI complementarity arm.

Within the control group (bottom panel of Table 3), men, older respondents, and university graduates more frequently regard public intervention as useful, whereas women and younger respondents more frequently report expected household financial loss and perceive a higher risk of

[15] Appendix D reports the English translations of the original questionnaire wording, while Panel B of Table D1 provides outcome definitions and descriptive statistics.

[16] The full wording of the petition and accompanying instructions is reported in Appendix D.

[17] Dechezleprêtre et al. (2025) similarly elicit anonymous support for a petition calling for climate action, informing respondents that the share supporting it will be communicated to their country's government. Although the authors describe this as willingness to sign a petition, their questionnaire explicitly states that no signature is required.

job loss. Regional differences in these perceptions are comparatively modest: perceived usefulness of public intervention and expected household financial loss are higher in the North, while perceived risk of job loss is highest in the South and Islands. Petition support varies little by gender and education but declines with age: 84.3% of respondents aged 18–34 support the petition, compared with 77.3% of those aged 35–49 and 76.5% of those aged 50–60. A geographical gap also emerges, with petition support at 82.8% in the South and Islands, 79.3% in the Center, and 75.9% in the North.

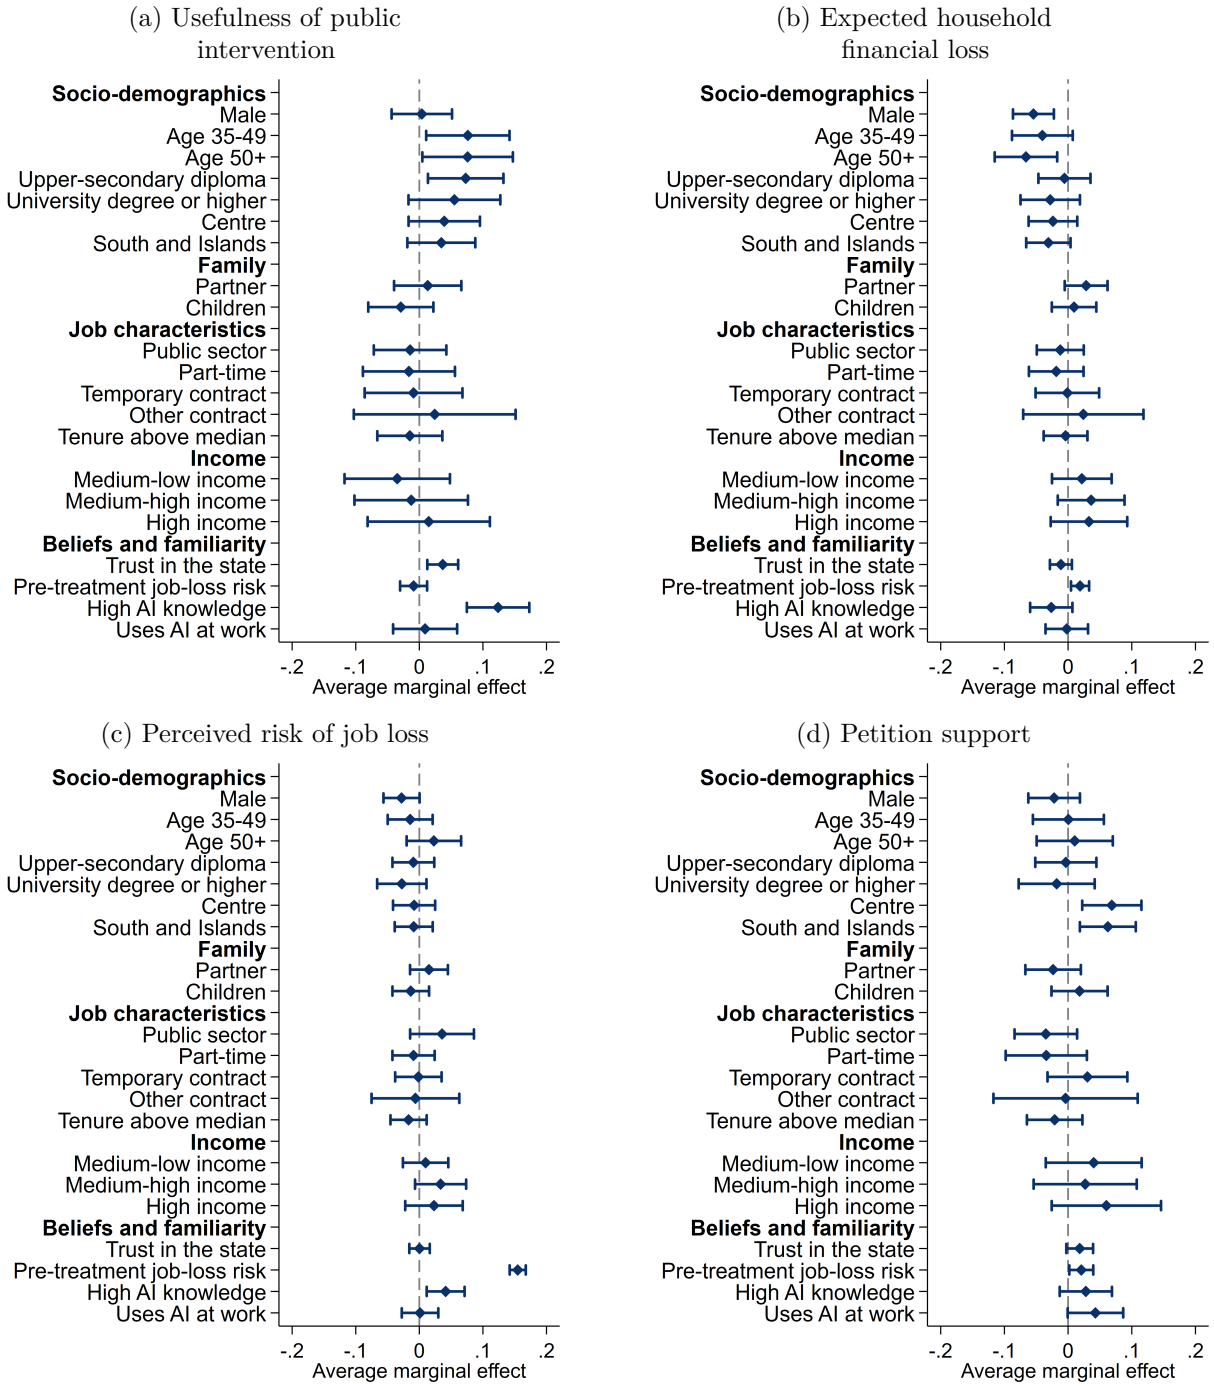
Table 3: Perceptions and real-stake behavior

Group	Category	N	Usefulness of public intervention	Expected household financial loss	Perceived risk of job loss	Petition support
All sample						
	Control	601	64.1	11.0	14.8	78.4
	AI substitution	502	63.5	14.9	15.9	79.3
	AI complementarity	500	65.8	10.6	13.2	80.6
	Reskilling	502	63.7	11.0	15.5	80.1
Control group						
Gender	Male	313	67.4	7.0	12.1	78.9
	Female	288	60.4	15.3	17.7	77.8
Age	18-34	115	59.1	16.5	19.1	84.3
	35-49	260	65.4	10.4	15.4	77.3
	50-60	226	65.0	8.8	11.9	76.5
Education	No university degree	438	62.8	11.2	14.6	78.3
	University degree	163	67.5	10.4	15.3	78.5
Region	North	315	65.1	12.4	14.0	75.9
	Center	135	63.0	8.1	14.8	79.3
	South and Islands	151	62.9	10.6	16.6	82.8

Notes: The table reports shares for 0 - 1 indicators for the whole sample by experimental arm and, within the control group, by sociodemographic group; reported values are multiplied by 100. Perceived usefulness of public intervention addressing AI's labor-market effects equals one for responses 4 or 5 on the original five-point scale; expected household financial loss equals one if the respondent expects such public intervention to leave their household financially somewhat or much worse off; perceived risk of job loss equals one if the respondent agrees or strongly agrees that they could lose their job within 36 months; and petition support equals one if the respondent is willing to support a petition on institutional action on AI. All indicators equal zero otherwise. Experimental-arm rows use the full main regression sample; control-group rows describe the 601 respondents assigned to the control group.

Figure 4 reports conditional associations between respondent characteristics and the four outcomes. All models control for treatment assignment and pre-treatment perceived risk of job loss. Perceived usefulness of public intervention is more clearly associated with trust in the state than with personal income: respondents with greater institutional trust are more likely to view intervention as useful, while differences across income groups are small and statistically insignificant. The association between trust and petition support is also positive, although less precisely estimated.

Figure 4: Analysis of individual perceptions and real-stake behavior



Notes: The figure reports average marginal effects from separate probit regressions for four binary outcomes. Perceived usefulness of public intervention addressing AI's labor-market effects equals one for responses 4 or 5 on the original five-point scale. Expected household financial loss equals one if the respondent expects such intervention to leave their household financially somewhat or much worse off. Perceived risk of job loss equals one for responses 4 or 5 on the five-point post-treatment perceived-risk scale, indicating agreement or strong agreement that the respondent could lose their job within the next 36 months. Petition support equals one if the respondent is willing to support the petition calling for institutional action on AI. All four models are estimated on the same 1,759 respondents. All specifications include sociodemographic, family, workplace, income, belief, and AI-familiarity covariates, pre-treatment perceived risk of job loss, and treatment-arm fixed effects. Horizontal bars report 95% confidence intervals based on heteroskedasticity-robust standard errors.

Geographical differences emerge for petition support: respondents in the Center and in the South and Islands are more likely to endorse the petition than those in the North.

Respondents who initially perceive a greater risk of losing their job are more likely to expect public intervention addressing AI's labor market effects to leave their household financially worse off. Yet they are also more likely to support the petition calling for institutional action. Thus, concerns about the potential household financial costs of intervention do not necessarily imply opposition to intervention itself.

Finally, respondents reporting high AI knowledge are more likely to regard public intervention as useful and to perceive a risk of job loss, even after controlling for initial risk perceptions. Workplace AI use is positively associated with petition support, although the evidence is weaker.

4.3 Average treatment effects on support for labor market policies

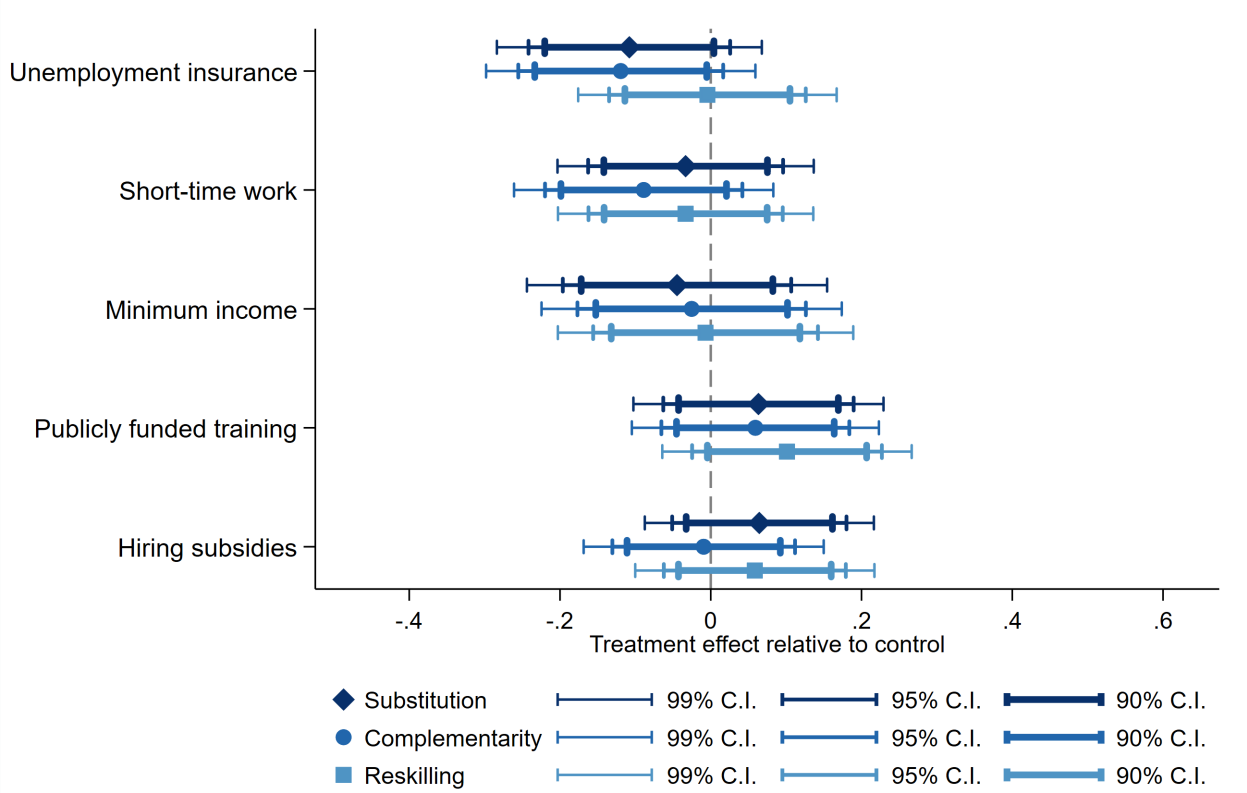
Figure 5 plots the average treatment effects of the three AI narratives on elicited support, measured on the original five-point scale, separately for each labor market policy, together with their confidence intervals. Table C1 in the Appendix reports the average treatment effects relative to the control group as estimated from Equation (1).

These estimates provide limited evidence of changes in average policy support. The AI complementarity narrative reduces support for unemployment insurance by 0.119 points on the five-point scale (which corresponds to a 3.2% reduction with respect to the control group), although this estimate is significant only at the 10% level. The substitution narrative yields a similar point estimate for this policy, while the reskilling narrative produces its largest positive estimate for publicly funded training; neither is statistically significant. These results are informative in their own right: salient qualitative information on the role of AI does not substantially shift average support for labor market policy in this sample.

To assess the robustness of our results, we examine whether the estimated treatment effects are sensitive to alternative specifications and sample restrictions. Table F1 in the Appendix reproduces in Panel A the main treatment-effect estimates (as reported in Table C1 and illustrated in Figure 5), while Panels B, C and D present three robustness checks. Overall, the results remain virtually unchanged. Panel B includes occupation fixed effects. Panel C estimates treatment effects conditional

on a broader set of covariates. Panel D excludes interviews with unusually short or long durations.[18]

Figure 5: Average treatment effects on labor market policies



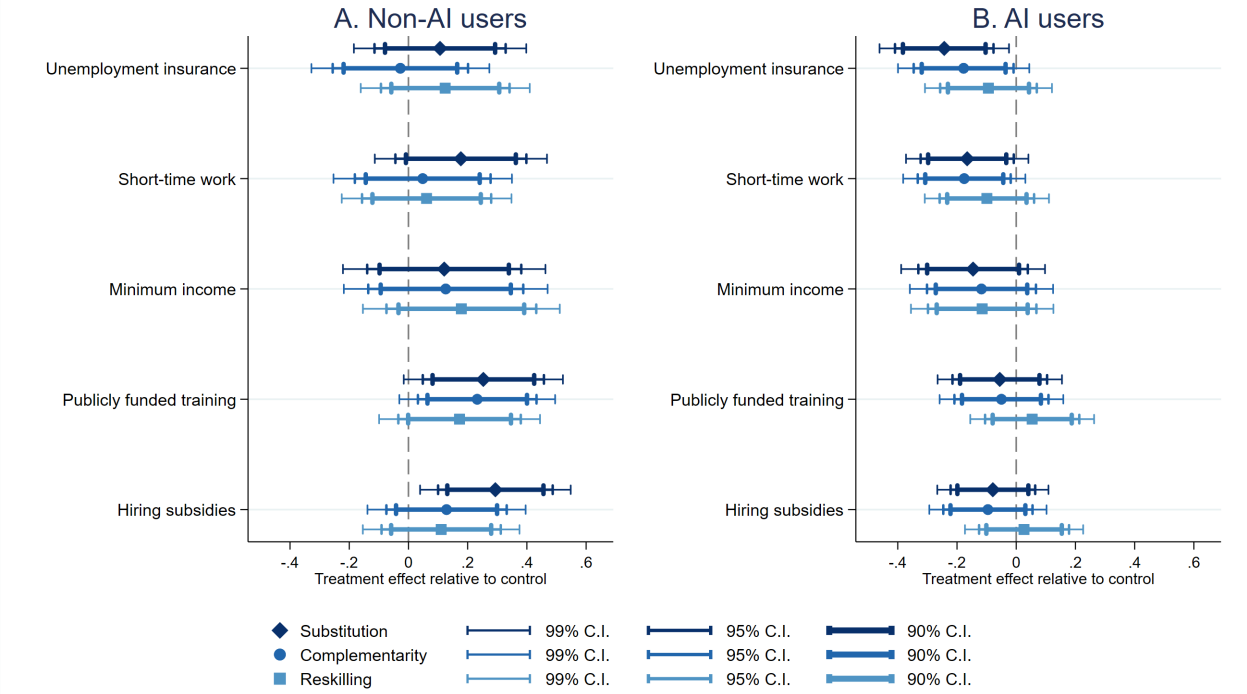
Notes: The figure reports treatment effects on the five policy-support outcomes, measured on their original five-point scales. Each point is the estimated difference in policy support between respondents assigned to the indicated treatment and respondents assigned to the control group. The estimates come from separate OLS models estimated on 2,105 respondents. Each model includes indicators for the three treatment arms; no additional covariates are included. Horizontal bars report 99%, 95%, and 90% confidence intervals based on heteroskedasticity-robust standard errors.

4.4 Effect heterogeneity by workplace AI use

Small average effects may nevertheless mask heterogeneous responses. We therefore investigate whether the effects of the AI narratives on policy preferences differ between workers who use AI for work and those who do not. Figure 6 presents the estimated treatment effects for non-users

[18] In Panel B, occupations are defined according to the one-digit Italian occupational classification. For the 309 respondents with missing occupation information, we introduce an additional category to preserve the baseline sample. The effect of AI complementarity on unemployment insurance is no longer significant at the 10% level. In Panel C, including the pre-treatment characteristics used in the balancing regressions implies using a restricted sample due to item-non response, and leads to confirming the effect of AI complementarity on unemployment insurance, while the effect of AI substitution also reaches significance at 10% level. For estimations reported in Panel D, within each experimental arm, we exclude interviews in the bottom and top 10% of the interview-duration distribution, measured in minutes using page timers, retaining 1,685 of the 2,105 respondents. The significance of the AI complementarity effect on unemployment insurance is confirmed (at the 5% level), while its effect on short-time work also reaches significance at the 10% level.

Figure 6: Treatment effects on policy support by workplace AI use



Notes: The figure reports treatment effects on the five policy-support outcomes, measured on their original five-point scales. Each point is the estimated difference in policy support between respondents assigned to the indicated treatment and respondents assigned to the control group within the same workplace AI-use group. Panel A reports these treatment-control differences among non-users, while Panel B reports them among AI users. The estimates come from separate OLS models estimated on 2,104 respondents: 820 non-users and 1,284 AI users. Each model includes treatment-arm indicators, an AI-user indicator, and their interactions; no additional covariates are included. Horizontal bars report 99%, 95%, and 90% confidence intervals based on heteroskedasticity-robust standard errors.

(Panel A) and AI users (Panel B), while Appendix Table C2 reports the corresponding estimates of Equation (2) for the five policy support outcomes. The effects for non-users correspond to the treatment coefficients, whereas those for AI users are obtained as linear combinations of the treatment coefficients and their interactions with workplace AI use.

The results reveal heterogeneous responses to the AI narratives depending on actual AI use for work. Among non-users (Panel A), both the substitution and complementarity treatments increase support for publicly funded training by 0.253 and 0.232 points, respectively. These effects correspond to approximately 6.5% and 6.0% of the non-users' control-group mean reported in Appendix Table C2. The substitution treatment also increases support for hiring subsidies by 0.294 points, equivalent to 7.5% of the corresponding control-group mean. None of the remaining treatment effects among non-users is statistically significant.

Turning to workers who use AI for work (Panel B), both narratives reduce support for unemploy-

ment insurance and short-time work schemes. The substitution treatment lowers support by 0.243 and 0.166 points, respectively, corresponding to reductions of around 6.5% and 4.6% relative to the control-group means for AI users. The corresponding effects of the complementarity treatment are -0.178 and -0.176 points, equivalent to reductions of 4.7% and 4.8%. Estimated effects on the minimum income schemes are also negative for both treatments, but are not statistically significant. Neither narrative produces a statistically significant change in support for training or hiring subsidies among AI users. The reskilling treatment yields no statistically significant effects in either group. Overall, the evidence suggests that treatment-effect heterogeneity is strongest for the substitution narrative, although it also extends to selected outcomes under the complementarity treatment. These findings remain robust to the same alternative specifications and sample restrictions considered for the average treatment effects.[19]

The results show that the treatment narratives influence workers' support for strengthening labor market policies, with responses differing by actual workplace AI use. This heterogeneity is particularly evident for the substitution narrative: among non-users, it increases support for strengthening publicly funded training and hiring subsidies; among AI users, it reduces support for strengthening unemployment insurance and short-time work schemes. The complementarity narrative produces a partially similar pattern, although evidence of differences between the two groups is less extensive.

5 Interpreting heterogeneous policy responses

Our findings show that limited average treatment effects conceal heterogeneous responses by workplace AI use. To help interpret how policy responses vary across narratives and between AI users and non-users, we examine treatment effects on perceived personal job loss risk, the perceived usefulness of public intervention for society, expected household financial consequences of intervention, and support for a petition calling for institutional action.

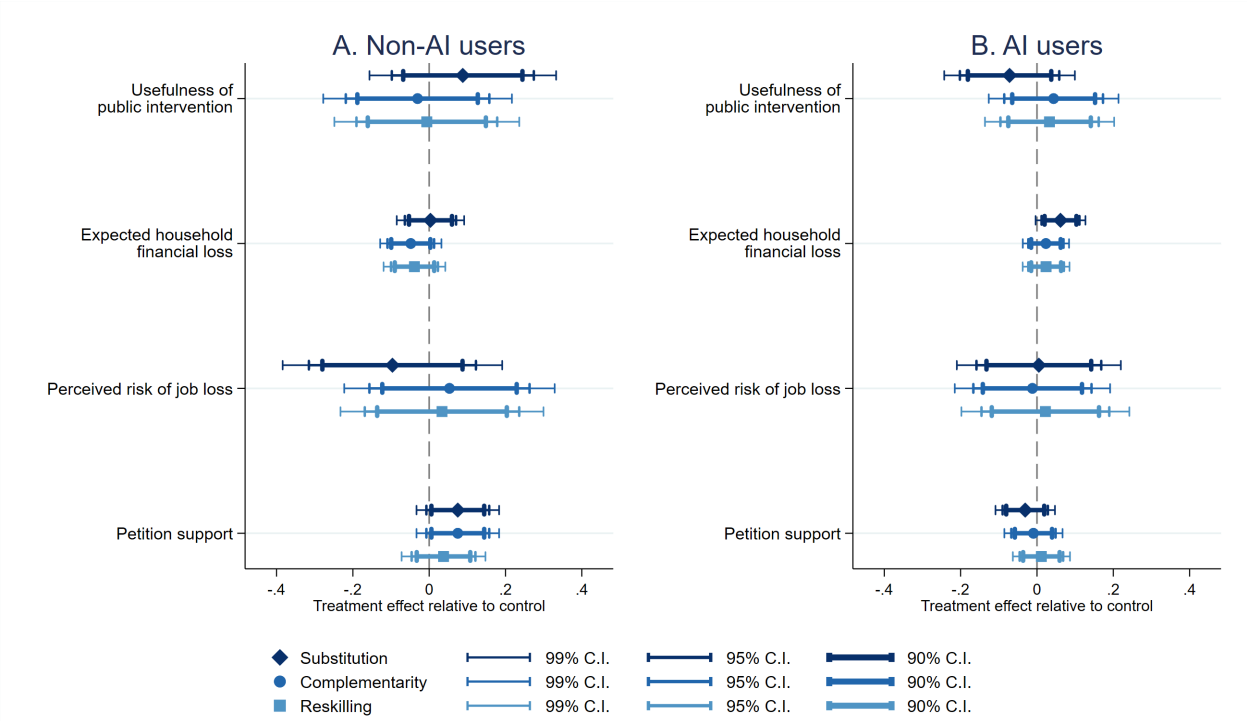
A natural starting point is the insurance motive for social protection. From this perspective, workers who perceive a greater labor market risk should place greater value on protection against future income losses. Evidence from the Great Recession is consistent with this mechanism: personal

[19] Tables F2 and F3 in the Appendix report the same set of robustness checks carried out for the average treatment effects in Table F1. The main patterns remain largely unchanged.

economic hardship, particularly job loss, increased support for welfare spending (Margalit, 2013). Higher local unemployment has also been associated with greater demand for unemployment benefits among employed workers, consistent with a demand for insurance against their own unemployment risk (Backus and Esteller-Moré, 2017). Similar patterns arise in the context of technological change, where workers who perceive greater technology-related employment risks tend to favor compensatory and protective policies (Busemeyer et al., 2023). The insurance motive, however, depends on workers relating broader labor market developments to their own employment prospects. Exposure to aggregate unemployment risk need not translate into greater demand for public intervention if workers are unaware of that risk or do not regard it as personally salient (Ahrens, 2024). Information about AI may therefore affect assessments of public intervention without necessarily changing perceived personal unemployment risk. Figure 7 reports the treatment effects separately for non-users and workplace AI users, while Table C3 in the Appendix provides the corresponding regression estimates. None of the narratives produces a statistically significant change in perceived job loss risk or perceived usefulness of public intervention in either group. Among non-users, both substitution and complementarity increase petition support by 7.5 percentage points, although these effects are significant only at the 10% level. Together with the increases in support for strengthening publicly funded training (see Figure 6), these results suggest that both narratives can encourage willingness to endorse a public response to technological change.

Among AI users, the substitution narrative increases by 6.2 percentage points the probability of expecting public intervention to leave their household financially worse off. Its estimated effect on perceived personal job loss risk is close to zero, and we detect no significant change in perceived policy usefulness or petition support. Neither complementarity nor reskilling produces statistically significant changes in these perceptions or in petition support among AI users. One possible interpretation of AI users' responses to the substitution treatment concerns the balance between the expected household financial impact and insurance benefits of intervention. Financing considerations may be particularly salient in our setting because, before evaluating the policy instruments, all respondents receive the same information about aggregate spending on labor market and social policies. Against this background, the substitution narrative may draw attention to the potential financing requirements associated with greater displacement. If AI users associate greater displacement with increased spending on unemployment insurance and short-time work

Figure 7: Treatment effects on perceptions and real-stakes behavior by workplace AI use



Notes: The figure reports treatment effects separately for respondents who do not use AI for work (Panel A) and workplace AI users (Panel B). Estimates come from OLS regressions that include indicators for the three treatment arms, an indicator for workplace AI use, and their interactions. The reported effects compare each treatment arm with the control group within the same AI-use category. Usefulness of public intervention in governing AI impact and perceived risk of job loss are measured on their original five-point scales, with higher values indicating greater perceived usefulness and greater perceived risk, respectively. Expected household financial loss equals one if the respondent expects public intervention to leave their household financially somewhat or much worse off. Petition support is binary as originally measured. Effects on perceived usefulness of public intervention in governing AI impact and perceived risk of job loss are therefore expressed in scale points, whereas effects on negative household impact and petition support are expressed as changes in probability. Models include no additional controls and use heteroskedasticity-robust standard errors. Horizontal intervals report 99, 95, and 90% confidence intervals.

without perceiving a greater likelihood of needing these programs themselves, strengthening them may appear less attractive. This interpretation is consistent with experimental evidence showing that making the tax costs of social policies salient can reduce support for unemployment benefits and redistribution (Naumann, 2018).

Complementarity may involve a different assessment of the benefits of intervention. Among non-users, information about AI's productive potential increases both support for strengthening publicly funded training (as shown in Figure 6), and willingness to support institutional action (as shown in Figure 7). One plausible interpretation is that the treatment expose workers to a novel information which makes the need to adapt to technological change more salient. Among AI users, by contrast, Figure 6 shows no comparable increase in support for training and a reduction in

support for strengthening unemployment insurance and short-time work. For these workers, the complementarity narrative may reinforce existing beliefs about AI's productive potential consistent with their direct workplace experience. Consistent with this interpretation Figure 7 also shows no comparable increase in willingness to support institutional action among AI users.

Overall, the results admit different interpretations of workers' responses to substitution and complementarity. Under substitution, the increase in expected household financial losses among AI users may reflect concerns about the costs of intervention relative to its perceived insurance benefits. Under complementarity, the differences between users and non-users may reflect how they assess the value of training and the perceived need for institutional support in light of AI's productive potential. These interpretations remain suggestive: we do not directly measure expectations about the costs of financing public intervention or the productivity gains from AI, and our analysis does not identify a causal pathway from the measured perceptions to support for strengthening particular policies.

6 Discussion and conclusion

This paper studies how workers translate alternative narratives about artificial intelligence into support for labor market policies. Using a four-arm survey experiment on a representative sample of Italian employees, we compare narratives emphasizing the AI's potential for substitution, complementarity, and reskilling. Average treatment effects are small, but they mask substantial heterogeneity by workplace AI use. Among non-users, the substitution narrative increases support for publicly funded training and hiring subsidies, while complementarity increases support for training. Among AI users, both narratives reduce support for unemployment insurance and short-time work schemes. The reskilling narrative produces no statistically significant effects on policy support in either group. These results suggest that workplace AI experience shapes not only the strength of workers' responses to information, but also the composition and direction of the resulting policy response.

Evidence on perceptions and petition support helps to interpret these findings. Neither perceived personal job-loss risk nor perceived usefulness of public intervention changes significantly under any narrative in either group. Among AI users, however, the substitution narrative increases the probability of expecting public intervention to leave their household financially worse off. Among

non-users, both substitution and complementarity increase petition support, although the evidence is weaker. These patterns suggest that responses to information about AI may involve assessments of intervention that extend beyond personal unemployment risk. For AI users, the substitution results are consistent with greater concern about the household financial costs of intervention without a detectable increase in perceived job-loss risk. However, we do not observe which financial costs respondents anticipate, nor do we establish that these perceptions cause changes in policy preferences. A different argument (reduced need for passive policies) explains the negative effects of complementarity among AI users. In this case, as expected, we do not find a statistically significant increase in the probability of expecting public intervention to leave their household financially worse off.

Our findings also raise questions about how policy demand might evolve as AI becomes more widespread and its labor market consequences more visible. Although our experiment does not forecast future preferences, the heterogeneity analysis may be informative about how workers respond as evidence of these consequences becomes more salient. In particular, among actual AI users, information emphasizing AI's substitution potential reduces support for strengthening protective passive labor policies. Information about a broader labor-market threat can therefore affect policy preferences before workers perceive a corresponding increase in their own employment risk.

Such responses need not persist as workers gain experience with AI and observe its consequences. If employment losses become more widespread or personally salient, workers may revise both their perceived risk and their willingness to strengthen protective policies. Recent evidence from Japan is consistent with the possibility that preferences evolve over time: as generative AI diffused, workers in more AI-exposed occupations moved toward greater support for redistribution relative to less exposed workers (Okubo and Wagner, 2026). Our estimates therefore do not imply a persistent relationship between workplace AI use and lower demand for social insurance.[20]

One possible implication is a policy-preparedness problem: public demand for stronger protection may emerge later than the need to design and finance it. Support for expanding social insurance may remain limited until employment risks become personally salient, even though preparing appropriate institutions may be easier before large-scale displacement materializes. If perceived vulnerability

[20] Our estimates do not identify the causal effect of AI use itself: they identify responses to randomized information within groups defined by their existing use of the technology.

subsequently increases, demand for compensation could rise when the scope for anticipatory action has already narrowed. Understanding how information and experience shape willingness to strengthen labor market policies is therefore relevant not only to which interventions attract support, but also to whether that support develops early enough to facilitate policy preparedness.

APPENDIX

A Italian labor market policies and their financing

We consider a set of five Italian policies that include unemployment protection schemes, income support and active labor market policies. These policies differ not only in their objectives and target groups, but also in how they are financed. These institutional differences are relevant for interpreting workers' preferences regarding whether or not to strengthen these policies in a scenario of widespread AI adoption in the labor market.

Protection against job loss takes the form of unemployment benefits and short-time work schemes. They are primarily insurance-based instruments delivered by the National Institute for Social Security (INPS), and are primarily financed through employment-based social-security contributions (collected by INPS), including additional contributions associated with fixed-term contracts, dismissals, or firms' use of short-time work schemes. The main unemployment benefit for dependent employees is NASpI (*Nuova Assicurazione Sociale per l'Impiego*), which provides income support following involuntary job loss. Short-time work support is provided through the *Cassa Integrazione Guadagni* system, which allows firms to reduce or temporarily suspend working hours while preserving the employment relationship. Contribution revenues, however, do not necessarily cover the full cost of these programs. When additional resources are required, financing may also come from transfers from the State budget to INPS. Their financing can therefore ultimately involve both payroll contributions and general taxation.

Minimum-income support has a different institutional basis. Since 2024, the main instrument has been the *Assegno di Inclusione*, complemented by the *Supporto per la Formazione e il Lavoro*. These are social-assistance programs financed through the general government budget rather than through dedicated unemployment-insurance contributions.

Active policies take the form of training programs for workers, which largely rely on national and European resources. Programs such as GOL (*Garanzia Occupabilità Lavoratori*) and the *Fondo Nuove Competenze* support reskilling, upskilling, and aim at facilitating labor market transitions. Employment and hiring subsidies usually involve reductions in, or exemptions from, employers' social-security contributions.

The five policy instruments therefore span different margins of labor market policy: insurance against income loss, employment preservation, redistribution, investment in workers' skills, and support for labor demand. Their different financing arrangements are relevant for how workers may evaluate an expansion of labor market intervention. Greater use of insurance-based programs may require additional resources within the contributory system and, when those resources are insufficient, from the general government budget. More generally, strengthening labor market policies may therefore imply additional financing needs borne through social contributions, general taxation, or both.

Our survey does not ask respondents which specific financing channel they have in mind when evaluating these policies, rather it elicit support to strengthening the specific selected policies described above. Nevertheless, the institutional structure makes it plausible that workers may perceive an expansion of labor market intervention as having financial consequences for their households.

B Descriptives

Table B1: Sample characteristics by workplace AI use

	Non-users (1)	AI users (2)	Difference (3)	p-value (4)
Panel A. Socio-demographic characteristics				
Age	46.623	42.599	-4.024	< 0.001
Female	0.482	0.444	-0.038	0.090
University degree or higher	0.150	0.357	0.207	< 0.001
North	0.563	0.512	-0.051	0.022
Center	0.182	0.227	0.046	0.010
South and Islands	0.255	0.260	0.005	0.788
Panel B. Household characteristics				
Household size	2.803	2.904	0.101	0.040
Lives with partner	0.670	0.700	0.031	0.141
Has children	0.607	0.561	-0.047	0.034
Number of children	0.961	0.886	-0.075	0.071
Panel C. Job characteristics				
Public sector	0.172	0.223	0.052	0.003
Part-time	0.259	0.162	-0.097	< 0.001
Permanent contract	0.850	0.834	-0.016	0.328
Temporary contract	0.113	0.124	0.010	0.469
Workplace has more than 15 employees	0.667	0.708	0.041	0.052
Days worked per week	5.162	5.152	-0.010	0.733
Hours worked per week	35.268	35.349	0.081	0.877
Current job tenure (years)	12.533	11.070	-1.463	0.003
Panel D. Income and financial situation				
Low personal income	0.189	0.115	-0.074	< 0.001
Medium-low personal income	0.304	0.241	-0.063	0.002
Medium-high personal income	0.366	0.395	0.029	0.191
High personal income	0.140	0.249	0.108	< 0.001
Make ends meet	3.116	3.375	0.259	< 0.001
Panel E. AI familiarity and pre-treatment belief				
High AI knowledge	0.247	0.644	0.398	< 0.001
Perceived risk of job loss before treatment	2.349	2.424	0.075	0.121
Observations	820	1,284		

Notes: The table reports sample means separately for non-users and workplace AI users. For binary variables, means represent sample shares. Workplace AI users report using AI for work every day, several times per week, or rarely; non-users report never using it. Workplace size equals one for workplaces with more than 15 employees. Personal income refers to net monthly income, grouped into four categories: low, medium-low, medium-high, and high. Column (3) reports the difference between AI users and non-users. Column (4) reports p-values from OLS regressions of each characteristic on an indicator for workplace AI use, using heteroskedasticity-robust standard errors.

Table B2: Treatment-arm balance among workplace AI users

	Control (1)	AI substitution (2)	AI complementarity (3)	Reskilling (4)	Joint-test p-value (5)
Socio-demographic characteristics					
Age	43.243	42.435	43.078	41.459	0.151
Female	0.441	0.406	0.477	0.453	0.348
University degree or higher	0.348	0.344	0.373	0.365	0.858
North	0.487	0.503	0.536	0.530	0.540
Center	0.257	0.218	0.212	0.216	0.481
South and Islands	0.257	0.279	0.252	0.253	0.857
Household characteristics					
Household size	2.949	2.752	3.016	2.890	0.014
Lives with partner	0.711	0.666	0.748	0.672	0.085
Has children	0.604	0.519	0.578	0.530	0.089
Number of children	0.976	0.799	0.935	0.811	0.040
Job characteristics					
Public sector	0.215	0.240	0.229	0.209	0.796
Part-time	0.158	0.140	0.186	0.166	0.471
Permanent contract	0.858	0.818	0.824	0.831	0.466
Temporary contract	0.091	0.146	0.137	0.128	0.092
Workplace has more than 15 employees	0.713	0.726	0.721	0.672	0.499
Days worked per week	5.111	5.167	5.153	5.186	0.344
Hours worked per week	35.102	35.142	35.075	36.161	0.595
Current job tenure (years)	11.336	10.475	11.021	11.410	0.667
Income and financial situation					
Low personal income	0.113	0.108	0.122	0.117	0.960
Medium-low personal income	0.229	0.212	0.233	0.294	0.130
Medium-high personal income	0.408	0.469	0.365	0.337	0.008
High personal income	0.249	0.212	0.280	0.252	0.284
Make ends meet	3.322	3.350	3.451	3.390	0.479
AI familiarity and pre-treatment belief					
High AI knowledge	0.644	0.636	0.645	0.652	0.984
Perceived risk of job loss before treatment	2.444	2.380	2.389	2.480	0.624
Observations	374	308	306	296	

Notes: The table reports sample means by experimental arm among workplace AI users. Workplace AI users are respondents who report using AI for work every day, several times per week, or rarely. For binary variables, means represent sample shares. Workplace size equals one for workplaces with more than 15 employees. Personal income refers to net monthly income, grouped into four categories: low, medium-low, medium-high, and high. The last column reports p-values for joint tests of the null hypothesis that the three treatment coefficients are jointly zero in OLS regressions of each characteristic on treatment assignment, using heteroskedasticity-robust standard errors. The control group is the omitted category.

Table B3: Treatment-arm balance among workplace AI non-users

	Control (1)	AI substitution (2)	AI complementarity (3)	Reskilling (4)	Joint-test p-value (5)
Socio-demographic characteristics					
Age	47.141	45.851	46.655	46.751	0.568
Female	0.542	0.423	0.536	0.420	0.009
University degree or higher	0.145	0.144	0.165	0.146	0.935
North	0.586	0.603	0.500	0.561	0.184
Center	0.172	0.149	0.237	0.171	0.165
South and Islands	0.242	0.247	0.263	0.268	0.918
Household characteristics					
Household size	2.831	2.840	2.813	2.725	0.711
Lives with partner	0.670	0.639	0.696	0.673	0.702
Has children	0.652	0.593	0.598	0.580	0.420
Number of children	1.053	0.923	0.933	0.922	0.395
Job characteristics					
Public sector	0.200	0.180	0.161	0.141	0.413
Part-time	0.260	0.278	0.299	0.200	0.104
Permanent contract	0.894	0.814	0.861	0.824	0.068
Temporary contract	0.075	0.139	0.098	0.146	0.053
Workplace has more than 15 employees	0.680	0.720	0.605	0.656	0.133
Days worked per week	5.164	5.187	5.203	5.095	0.354
Hours worked per week	35.514	34.963	33.965	36.533	0.153
Current job tenure (years)	14.158	11.916	11.702	12.068	0.091
Income and financial situation					
Low personal income	0.173	0.189	0.268	0.130	0.010
Medium-low personal income	0.276	0.322	0.311	0.314	0.741
Medium-high personal income	0.421	0.339	0.284	0.411	0.014
High personal income	0.131	0.150	0.137	0.146	0.947
Make ends meet	3.141	3.068	2.979	3.264	0.072
AI familiarity and pre-treatment belief					
High AI knowledge	0.221	0.237	0.263	0.268	0.641
Perceived risk of job loss before treatment	2.308	2.304	2.407	2.380	0.715
Observations	227	194	194	205	

Notes: The table reports sample means by experimental arm among non-users. Non-users are respondents who report never using AI for work. For binary variables, means represent sample shares. Workplace size equals one for workplaces with more than 15 employees. Personal income refers to net monthly income, grouped into four categories: low, medium-low, medium-high, and high. The last column reports p-values for joint tests of the null hypothesis that the three treatment coefficients are jointly zero in OLS regressions of each characteristic on treatment assignment, using heteroskedasticity-robust standard errors. The control group is the omitted category.

C Regression results

This appendix reports the regression estimates underlying the treatment-effect figures in the main text. Table C1 reports average treatment effects on policy support, Table C2 reports treatment interactions with workplace AI use for the policy outcomes, and Table C3 reports the corresponding interactions for the additional outcomes.

C.1 Average treatment effects on policy support

Table C1: Average Treatment effects on support for labor market policies

	Support for Labour Market Policies				
	Unemp. insurance	Short-time work	Minimum income	Publicly funded training	Hiring subsidies
	(1)	(2)	(3)	(4)	(5)
AI Substitution	-0.108	-0.033	-0.045	0.063	0.064
	(0.068)	(0.066)	(0.077)	(0.064)	(0.059)
AI Complementarity	-0.119*	-0.089	-0.025	0.059	-0.009
	(0.069)	(0.067)	(0.077)	(0.064)	(0.062)
Reskilling	-0.004	-0.033	-0.007	0.101	0.058
	(0.067)	(0.066)	(0.076)	(0.064)	(0.062)
Control mean	3.755	3.629	3.399	3.907	3.983
Observations	2,105	2,105	2,105	2,105	2,105
Adj. R-squared	0.001	-0.001	-0.001	-0.000	-0.000

Notes: Each column reports OLS estimates from a model without additional controls. Outcome variables are measured on a five-point scale. The control group is the reference category. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

C.2 Policy-support effects by workplace AI use

Table C2: Average treatment effects on support for labor market policies - workplace AI use heterogeneity

	Labour Market Policies				
	Unemp. insurance	Short-time work	Minimum income	Publicly funded training	Hiring subsidies
	(1)	(2)	(3)	(4)	(5)
AI substitution	0.107	0.177	0.121	0.253**	0.294***
	(0.113)	(0.113)	(0.133)	(0.104)	(0.099)
AI complementarity	-0.027	0.048	0.126	0.232**	0.129
	(0.117)	(0.117)	(0.134)	(0.102)	(0.104)
Reskilling	0.124	0.061	0.179	0.173	0.110
	(0.111)	(0.111)	(0.129)	(0.106)	(0.103)
AI user	0.010	0.020	0.231**	0.041	0.072
	(0.097)	(0.095)	(0.112)	(0.095)	(0.090)
AI substitution × AI user	-0.350**	-0.343**	-0.267	-0.309**	-0.373***
	(0.141)	(0.138)	(0.163)	(0.132)	(0.123)
AI complementarity × AI user	-0.150	-0.224	-0.243	-0.282**	-0.224*
	(0.145)	(0.142)	(0.163)	(0.130)	(0.129)
Reskilling × AI user	-0.218	-0.161	-0.294*	-0.119	-0.084
	(0.139)	(0.138)	(0.159)	(0.133)	(0.129)
Control mean, non-AI users	3.749	3.617	3.256	3.881	3.938
Control mean, AI users	3.759	3.636	3.487	3.922	4.011
Observations	2,104	2,104	2,104	2,104	2,104
Adjusted R-squared	0.007	0.005	-0.001	0.005	0.005
Joint p-value: treatment × AI user	0.095	0.094	0.228	0.064	0.014

Notes: The table reports OLS estimates of the outcomes listed in the columns on treatment indicators, an indicator for AI use for work, and their interactions. All outcomes are measured on a five-point scale, ranging from 1 ("not at all in favor") to 5 ("strongly in favor"). The omitted categories are the control group and non-AI users. Treatment coefficients refer to non-AI users, while interaction terms capture differences in treatment effects between AI users and non-users. The joint p-value tests whether the three interaction coefficients are jointly zero. No additional controls are included. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

C.3 Perceptions and real-stake behavior - workplace AI use heterogeneity

Table C3: Average Treatment Effects on Perceptions and real-stake behavior - workplace AI use heterogeneity

	Usefulness of public intervention (1)	Expected household financial loss (2)	Petition support (3)	Perceived risk of job loss (4)
AI substitution	0.088	0.003	0.075*	-0.096
	(0.095)	(0.034)	(0.042)	(0.112)
AI complementarity	-0.030	-0.048	0.075*	0.053
	(0.096)	(0.031)	(0.042)	(0.107)
Reskilling	-0.006	-0.039	0.038	0.034
	(0.094)	(0.031)	(0.043)	(0.103)
AI user	0.071	-0.050*	0.113***	0.104
	(0.078)	(0.028)	(0.036)	(0.092)
AI substitution × AI user	-0.160	0.058	-0.106**	0.101
	(0.116)	(0.043)	(0.052)	(0.139)
AI complementarity × AI user	0.074	0.072*	-0.084	-0.065
	(0.116)	(0.039)	(0.051)	(0.133)
Reskilling × AI user	0.039	0.062	-0.026	-0.011
	(0.115)	(0.039)	(0.052)	(0.134)
Control mean, non-AI users	3.670	0.141	0.714	2.313
Control mean, AI users	3.741	0.091	0.826	2.417
Observations	2,104	2,104	2,104	2,104
Adjusted R-squared	0.000	0.002	0.005	0.000
Joint p-value: treatment × AI user	0.224	0.248	0.143	0.706

Notes: The table reports OLS estimates for the additional outcomes on treatment indicators, an indicator for AI use for work, and their interactions. Usefulness of public intervention in governing AI impact and perceived risk of job loss are measured on their original five-point scales, with higher values indicating greater usefulness and greater perceived risk, respectively. Expected household financial loss equals one if the respondent expects public intervention to leave the household financially somewhat or much worse off. Petition support equals one if the respondent is willing to support the petition calling for institutional action on AI. The omitted categories are the control group and non-AI users. Treatment coefficients refer to non-AI users, while interaction terms capture differences in treatment effects between AI users and non-users. The joint p-value tests whether the three interaction coefficients are jointly zero. No additional controls are included. One respondent did not report AI use for work and is excluded. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

D Questionnaire materials

The questionnaire was fielded in Italian. This appendix reports close English translations of the treatment texts and of the key outcome wording used in the analysis. These details matter for interpretation because the active arms differ in relatively focused claims about AI, and because several post-treatment outcomes explicitly reference public intervention in governing AI's labor market impact.

D.1 Common introduction to the AI block

"We will now ask you some questions about Artificial Intelligence (AI)."

"By Artificial Intelligence (AI), we mean a set of technologies and computer programs that, by analyzing large quantities of data, can carry out tasks that would usually require human intelligence. These tasks include, for example, understanding and generating texts in different languages, recognizing and classifying images or sounds, making predictions, making decisions, or proposing solutions to problems."

D.2 Treatment wording

T1: Substitution frame.

"A recent scientific analysis of the exposure of the Italian labor market to Artificial Intelligence (AI) indicates that a distinctive feature of AI-based technologies is their ability to perform activities that were previously considered intrinsically human and difficult to automate. This means that among workers exposed to AI there are more people with medium-to-high levels of education. In particular, AI's ability to substitute for human labor appears to be greater in occupations that require specialized but standardized procedures."

T2: Complementarity frame.

"A recent scientific analysis of the exposure of the Italian labor market to Artificial Intelligence (AI) indicates that a distinctive feature of AI-based technologies is their ability to perform activities that were previously considered intrinsically human and difficult to automate. This means that among workers exposed to AI there are more people with medium-to-high levels of education. In particular, AI's ability to accompany and enhance human labor is generally greater in occupations that often involve complex, high-level decisions."

T3: Reskilling frame.

> "A recent scientific analysis indicates that workers who have completed at least one training course show a significantly higher impact of Artificial Intelligence than those who did not participate in courses. This finding highlights how continuous training and professional updating are key factors for the adoption and integration of AI technologies for work."

D.3 Key outcome wording

Labor market policies.

> "In light of recent and upcoming changes in the labor market linked to the use of Artificial Intelligence, how favorable would you be to strengthening each of the following public interventions? Please use a score from 1 = not at all favorable to 5 = very favorable. With the intermediate scores you can grade your judgment."
>
> Items shown in randomized order: unemployment insurance (NASpI / DIS-COLL); short-time work schemes and solidarity contracts; minimum-income policies such as the Inclusion Allowance; state-funded training courses for workers; tax incentives to firms to encourage new hiring.

Usefulness of public intervention.

> "How useful do you think public intervention is, for the well-being of society, in governing the impact of Artificial Intelligence on the labor market, for example through income-support measures, training courses for workers, and tax incentives to firms to encourage new hiring?"

Expected household financial impact.

> "Thinking about public interventions to govern the impact of Artificial Intelligence on the labor market (for example, income-support measures, training courses for workers, and tax incentives to firms to encourage new hiring), do you think that, from a financial point of view, your family would..." [lose a lot / lose somewhat / neither lose nor benefit / benefit somewhat / benefit a lot].

Perceived risk of job loss.

> "Do you think you could lose your job in the next 36 months?"

Petition support.

Respondents first read: "We now ask whether you would be willing to support the following petition, which calls for institutional intervention to govern the impact of Artificial Intelligence on firms and workers. The result of the petition, namely the share of participants in favor, will be communicated to the Observatory on AI adoption in the labor market (Ministry of Labor and Social Policies). The response remains anonymous and no signature is required."

Petition text: "I believe it is essential that institutions intervene urgently to govern the impact of Artificial Intelligence on the labor market. Clear rules are needed to protect employment, promote skills retraining, and ensure a fair transition toward an innovative economy. I call for concrete measures and direct action aimed at maximizing the benefits and minimizing the social risks of AI."

"Would you support this petition?"

Table D1: Outcome definitions and descriptive statistics

Outcome	Scale	Definition	N	Mean	SD	Min	Max
Panel A. Policy-support outcomes							
Unemployment insurance	1–5	Support for strengthening unemployment-insurance schemes	2,105	3.700	1.129	1.000	5.000
Short-time work schemes	1–5	Support for strengthening short-time work schemes	2,105	3.592	1.099	1.000	5.000
Minimum income	1–5	Support for strengthening the minimum income	2,105	3.381	1.268	1.000	5.000
Publicly funded training	1–5	Support for strengthening publicly funded training and reskilling	2,105	3.960	1.045	1.000	5.000
Hiring subsidies	1–5	Support for strengthening subsidies to firms that hire workers	2,105	4.010	1.002	1.000	5.000
Panel B. Perceptions and real-stake behavior							
Usefulness of public intervention	1–5	Perceived usefulness of public intervention	2,105	3.718	0.905	1.000	5.000
Expected household financial loss	0/1	Equals one if the respondent expects to be somewhat or much worse off	2,105	0.118	0.323	0.000	1.000
Perceived risk of job loss	1–5	Agreement that the respondent could lose their job within 36 months	2,105	2.378	1.082	1.000	5.000
Petition support	0/1	Equals one if the respondent is willing to support the petition	2,105	0.795	0.404	0.000	1.000
Panel C. Pre-treatment measure							
Pre-treatment perceived risk of job loss	1–5	Perceived-risk measure collected before treatment assignment	2,105	2.395	1.077	1.000	5.000

Notes: The table reports definitions and descriptive statistics for the outcomes used in the main analysis. Statistics refer to the main estimation sample. Policy support is measured from 1 (not at all in favor) to 5 (strongly in favor). Usefulness of public intervention in governing AI impact is measured from 1 (not at all useful) to 5 (very useful). Perceived risk of job loss uses a five-point agreement scale, with higher values indicating greater perceived risk. Usefulness of public intervention in governing AI impact and perceived risk of job loss enter the treatment-effect specifications on their original five-point scales. Expected household financial loss equals one if the respondent expects public intervention to leave the household financially somewhat or much worse off. Petition support equals one if the respondent is willing to support the petition. In the binary descriptive analyses, Usefulness of public intervention in governing AI impact and perceived risk of job loss are recoded as indicators equal to one for responses 4 or 5. The pre-treatment perceived-risk measure is used as a covariate rather than as a post-treatment outcome.

E Randomization

Table E1: Balance Analysis

	Substitution	Complementarity	Reskilling	Control
Male	0.047**	-0.040*	0.016	-0.023
	(0.021)	(0.021)	(0.021)	(0.022)
Age	-0.001	0.001	-0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Upper-secondary diploma	0.025	-0.043*	0.004	0.015
	(0.025)	(0.026)	(0.025)	(0.027)
University degree or higher	0.006	-0.008	0.007	-0.005
	(0.030)	(0.031)	(0.031)	(0.032)
Center	-0.023	0.021	-0.026	0.028
	(0.024)	(0.025)	(0.024)	(0.027)
South and Islands	-0.002	0.001	-0.012	0.013
	(0.024)	(0.024)	(0.023)	(0.025)
Lives with partner	-0.030	0.056**	-0.009	-0.016
	(0.023)	(0.022)	(0.024)	(0.025)
Has children	-0.013	-0.019	-0.014	0.046*
	(0.022)	(0.022)	(0.023)	(0.024)
Public sector	0.036	-0.007	-0.032	0.004
	(0.025)	(0.025)	(0.024)	(0.026)
Part-time	0.014	0.020	-0.030	-0.004
	(0.032)	(0.030)	(0.031)	(0.033)
Temporary contract	0.047	-0.014	0.033	-0.066**
	(0.034)	(0.032)	(0.033)	(0.031)
Other/nonstandard contract	0.018	-0.006	-0.043	0.031
	(0.054)	(0.054)	(0.048)	(0.056)
Medium-low income	-0.000	-0.040	0.048	-0.008
	(0.037)	(0.036)	(0.035)	(0.038)
Medium-high income	0.034	-0.065*	0.003	0.028
	(0.039)	(0.037)	(0.038)	(0.040)
High income	-0.008	-0.013	0.020	0.001
	(0.043)	(0.041)	(0.041)	(0.043)
AI user	-0.006	0.010	-0.019	0.015
	(0.021)	(0.021)	(0.021)	(0.022)
Constant	0.231***	0.262***	0.304***	0.202***
	(0.064)	(0.065)	(0.064)	(0.066)
Observations	1977	1977	1977	1977
Adjusted R-squared	0.002	0.005	0.000	0.001

Notes: The table reports OLS estimates from separate regressions of experimental-arm assignment indicators on respondent characteristics. Each column corresponds to a different experimental arm. The estimation sample includes 1,977 respondents with complete information on all covariates. Education is classified into three categories: lower secondary or less, upper-secondary diploma, and university degree or higher. The omitted categories are female, lower secondary or less, North, no partner, no children, private-sector employment, full-time employment, permanent contract, low income, and non-AI users. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F Robustness checks

F.1 Robustness of average treatment effects

We conduct several robustness checks. First, we assess whether the average treatment effects are sensitive to interview duration, conditioning on observed respondent characteristics, and occupation fixed effects. We then restrict the sample to the central 80% of the interview-duration distribution within each experimental arm, include the pre-treatment characteristics used in the balancing regressions, and separately add one-digit occupation fixed effects.[21] Table F1 reproduces in Panel A the main treatment-effect estimates (as reported in Table C1), while Panels B, C and D report estimates with occupation fixed effects, additional covariates, and interview-duration restrictions, respectively.

[21] To preserve the same estimation sample as the main regression (Panel A specification in Table F1), we retain all 2,105 respondents. We observe nine one-digit occupation categories, while the 309 respondents with missing occupation information are pooled into one additional category.

Table F1: Average treatment effects on support for labor market policies - robustness checks

	Unemployment insurance	Short-time work	Minimum income	Publicly funded training	Hiring subsidies
Panel A. Main, unconditional					
AI substitution	-0.108	-0.033	-0.045	0.063	0.064
	(0.068)	(0.066)	(0.077)	(0.064)	(0.059)
AI complementarity	-0.119*	-0.089	-0.025	0.059	-0.009
	(0.069)	(0.067)	(0.077)	(0.064)	(0.062)
Reskilling	-0.004	-0.033	-0.007	0.101	0.058
	(0.067)	(0.066)	(0.076)	(0.064)	(0.062)
Control mean	3.755	3.629	3.399	3.907	3.983
Observations	2,105	2,105	2,105	2,105	2,105
Adjusted R-squared	0.001	-0.001	-0.001	-0.000	-0.000
Panel B. Unconditional with occupation fixed effects					
AI substitution	-0.101	-0.026	-0.037	0.063	0.066
	(0.068)	(0.066)	(0.077)	(0.064)	(0.058)
AI complementarity	-0.113	-0.087	-0.025	0.055	-0.004
	(0.069)	(0.066)	(0.077)	(0.063)	(0.061)
Reskilling	0.001	-0.032	0.002	0.102	0.060
	(0.066)	(0.066)	(0.076)	(0.064)	(0.062)
Control mean	3.755	3.629	3.399	3.907	3.983
Observations	2,105	2,105	2,105	2,105	2,105
Adjusted R-squared	0.022	0.015	0.011	0.012	0.013
Panel C. Conditional					
AI substitution	-0.117*	-0.035	-0.096	0.077	0.066
	(0.070)	(0.069)	(0.080)	(0.067)	(0.061)
AI complementarity	-0.160**	-0.103	-0.061	0.061	-0.025
	(0.071)	(0.069)	(0.079)	(0.066)	(0.063)
Reskilling	-0.020	-0.066	-0.037	0.102	0.043
	(0.069)	(0.069)	(0.079)	(0.067)	(0.065)
Control mean	3.765	3.642	3.425	3.917	4.000
Observations	1,977	1,977	1,977	1,977	1,977
Adjusted R-squared	0.038	0.012	0.018	0.015	0.016
Panel D. Central 80 percent within treatment arm					
AI substitution	-0.088	-0.042	-0.124	0.071	0.074
	(0.076)	(0.074)	(0.087)	(0.071)	(0.066)
AI complementarity	-0.159**	-0.139*	-0.088	0.048	-0.006
	(0.080)	(0.076)	(0.088)	(0.071)	(0.070)
Reskilling	-0.003	-0.007	-0.020	0.101	0.099
	(0.075)	(0.073)	(0.086)	(0.072)	(0.069)
Control mean	3.769	3.649	3.418	3.917	3.983
Observations	1,685	1,685	1,685	1,685	1,685
Adjusted R-squared	0.002	0.001	-0.000	-0.000	0.000

Notes: The table reports OLS treatment effects on support for strengthening the five labor market policies, measured on their original five-point scales (1 = "not at all in favor"; 5 = "strongly in favor"). Each coefficient represents the effect of the indicated treatment relative to the control group. Panel A reproduces the main treatment-effect estimates reported in Table C1 (N = 2,105). Panel B adds one-digit occupation fixed effects, retaining nine observed categories and an additional category for the 309 respondents with missing occupation information, thereby preserving the baseline sample. Panel C adds individual controls for gender, age, education, geographic area, partner and parental status, public-sector employment, part-time status, contract type, income group, and workplace AI use. The estimation sample includes 1,977 respondents with non-missing information on all covariates. Panel D re-estimates the Panel A specification after excluding the shortest and longest 10% of interviews within each experimental arm, retaining 1,685 respondents. Interview duration is measured using page timers. Control means refer to the control group within the corresponding estimation sample. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F.2 Robustness of treatment effects by workplace AI use

We assess the robustness of treatment effects by workplace AI use using the same alternative specifications and sample restrictions considered for the average treatment effects. Specifically, we examine whether the estimated effects are sensitive to the inclusion of occupation fixed effects, conditioning on observed respondent characteristics, and restricting the sample to the central 80% of interview duration within each experimental arm. Table F2 reproduces in Panel A the main treatment effects by workplace AI use (as illustrated in Figure 6 and derived from Table C2) and reports estimates with occupation fixed effects in Panel B. Table F3 reports conditional treatment effects in Panel C and estimates with interview-duration restrictions in Panel D. Overall, the main patterns of treatment-effect heterogeneity remain largely unchanged across specifications.

Table F2: Treatment effects by workplace AI use: robustness checks (1)

	Unemployment insurance	Short-time work	Minimum income	Publicly funded training	Hiring subsidies
Panel A. Main, unconditional					
AI substitution					
Non-users	0.107 (0.113)	0.177 (0.113)	0.121 (0.133)	0.253** (0.104)	0.294*** (0.099)
AI users	-0.243*** (0.085)	-0.166** (0.080)	-0.146 (0.094)	-0.056 (0.082)	-0.079 (0.073)
Difference p-value	0.013	0.013	0.102	0.020	0.002
AI complementarity					
Non-users	-0.027 (0.117)	0.048 (0.117)	0.126 (0.134)	0.232** (0.102)	0.129 (0.104)
AI users	-0.178** (0.086)	-0.176** (0.080)	-0.117 (0.094)	-0.050 (0.081)	-0.096 (0.077)
Difference p-value	0.300	0.115	0.136	0.031	0.082
Reskilling					
Non-users	0.124 (0.111)	0.061 (0.111)	0.179 (0.129)	0.173 (0.106)	0.110 (0.103)
AI users	-0.094 (0.083)	-0.099 (0.081)	-0.115 (0.093)	0.054 (0.081)	0.026 (0.078)
Difference p-value	0.116	0.244	0.065	0.373	0.514
Control mean, non-users	3.749	3.617	3.256	3.881	3.938
Control mean, AI users	3.759	3.636	3.487	3.922	4.011
Observations, non-users	820	820	820	820	820
Observations, AI users	1,284	1,284	1,284	1,284	1,284
Panel B. Unconditional with occupation fixed effects					
AI substitution					
Non-users	0.124 (0.113)	0.202* (0.112)	0.142 (0.132)	0.268** (0.105)	0.313*** (0.098)
AI users	-0.240*** (0.084)	-0.167** (0.080)	-0.146 (0.094)	-0.065 (0.080)	-0.087 (0.072)
Difference p-value	0.010	0.008	0.077	0.012	0.001
AI complementarity					
Non-users	-0.009 (0.116)	0.069 (0.115)	0.143 (0.132)	0.256** (0.102)	0.154 (0.102)
AI users	-0.176** (0.085)	-0.184** (0.080)	-0.128 (0.093)	-0.070 (0.081)	-0.101 (0.076)
Difference p-value	0.243	0.071	0.093	0.012	0.046
Reskilling					
Non-users	0.142 (0.111)	0.078 (0.111)	0.206 (0.129)	0.186* (0.106)	0.120 (0.102)
AI users	-0.092 (0.083)	-0.104 (0.081)	-0.118 (0.093)	0.050 (0.081)	0.026 (0.078)
Difference p-value	0.091	0.187	0.041	0.307	0.462
Control mean, non-users	3.749	3.617	3.256	3.881	3.938
Control mean, AI users	3.759	3.636	3.487	3.922	4.011
Observations, non-users	820	820	820	820	820
Observations, AI users	1,284	1,284	1,284	1,284	1,284

Notes: The table reports OLS treatment effects on support for strengthening the five labor market policies, measured on their original five-point scales (1 = "not at all in favor"; 5 = "strongly in favor"). Treatment effects are reported separately for workplace AI users and non-users, relative to the control group. Panel A reproduces the main treatment effects by workplace AI use illustrated in Figure 6 and derived from the estimates reported in Table C2. Panel B adds one-digit occupation fixed effects, retaining nine observed categories and an additional category for the 308 respondents with missing occupation information, thereby preserving the Panel A sample. Effects for AI users are calculated as linear combinations of the treatment and interaction coefficients. Difference p-values test equality of treatment effects between AI users and non-users for each treatment. Workplace AI users report using AI for work every day, several times per week, or rarely; non-users report never using it. Control means and observation counts refer to the corresponding AI-use group and estimation sample. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F3: Treatment effects by workplace AI use: robustness checks (2)

	Unemployment insurance	Short-time work	Minimum income	Publicly funded training	Hiring subsidies
Panel C. Conditional					
AI substitution					
Non-users	0.084 (0.117)	0.160 (0.119)	0.031 (0.137)	0.270** (0.110)	0.335*** (0.103)
AI users	-0.241*** (0.087)	-0.155* (0.083)	-0.173* (0.097)	-0.042 (0.084)	-0.100 (0.075)
Difference p-value	0.026	0.030	0.224	0.024	<0.001
AI complementarity					
Non-users	-0.091 (0.121)	0.029 (0.124)	0.060 (0.138)	0.230** (0.107)	0.130 (0.108)
AI users	-0.201** (0.087)	-0.182** (0.082)	-0.134 (0.096)	-0.040 (0.083)	-0.117 (0.078)
Difference p-value	0.460	0.156	0.249	0.046	0.063
Reskilling					
Non-users	0.103 (0.115)	0.016 (0.117)	0.129 (0.134)	0.153 (0.112)	0.111 (0.109)
AI users	-0.095 (0.085)	-0.115 (0.085)	-0.141 (0.096)	0.073 (0.084)	0.004 (0.080)
Difference p-value	0.161	0.364	0.102	0.567	0.427
Control mean, non-users	3.746	3.620	3.282	3.897	3.930
Control mean, AI users	3.776	3.656	3.511	3.929	4.043
Observations, non-users	759	759	759	759	759
Observations, AI users	1,218	1,218	1,218	1,218	1,218
Panel D. Central 80 percent within treatment arm					
AI substitution					
Non-users	0.180 (0.129)	0.172 (0.125)	0.058 (0.148)	0.273** (0.117)	0.296*** (0.113)
AI users	-0.261*** (0.094)	-0.183** (0.090)	-0.236** (0.106)	-0.062 (0.090)	-0.069 (0.081)
Difference p-value	0.006	0.021	0.107	0.023	0.009
AI complementarity					
Non-users	0.024 (0.133)	-0.001 (0.130)	0.113 (0.147)	0.191* (0.114)	0.164 (0.116)
AI users	-0.279*** (0.100)	-0.232** (0.092)	-0.213** (0.108)	-0.047 (0.092)	-0.117 (0.088)
Difference p-value	0.068	0.146	0.074	0.102	0.054
Reskilling					
Non-users	0.221* (0.124)	0.072 (0.121)	0.212 (0.141)	0.159 (0.117)	0.149 (0.113)
AI users	-0.165* (0.095)	-0.066 (0.092)	-0.161 (0.108)	0.064 (0.092)	0.076 (0.087)
Difference p-value	0.014	0.363	0.037	0.520	0.610
Control mean, non-users	3.678	3.623	3.240	3.863	3.891
Control mean, AI users	3.826	3.664	3.527	3.950	4.040
Observations, non-users	685	685	685	685	685
Observations, AI users	999	999	999	999	999

Notes: The table reports OLS treatment effects on support for strengthening the five labor market policies, measured on their original five-point scales (1 = "not at all in favor"; 5 = "strongly in favor"). Treatment effects are reported separately for workplace AI users and non-users, relative to the control group. Panel C reports conditional estimates controlling for gender, age, education, geographic area, partner and parental status, public-sector employment, part-time status, contract type, and income group. The estimation sample includes 1,977 respondents with non-missing information on all covariates. Panel D re-estimates the Panel A specification of Table F2 after excluding the shortest and longest 10% of interviews within each experimental arm, retaining 1,684 respondents. Interview duration is measured as the sum of page timers. Trimming is performed within experimental arms, not separately by AI-use group. Effects for AI users are calculated as linear combinations of the treatment and interaction coefficients. Difference p-values test equality of treatment effects between AI users and non-users for each treatment. Workplace AI users report using AI for work every day, several times per week, or rarely; non-users report never using it. Control means and observation counts refer to the corresponding AI-use group and estimation sample. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this work the authors used ChatGPT to proofread and streamline the paper. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. Journal of Labor Economics, 40(S1):S293–S340.

Acemoglu, D., Gitmez, A. A., and Shadmehr, M. (2026). Automation and repression. NBER Working Paper No. 35336, National Bureau of Economic Research.

Acemoglu, D. and Restrepo, P. (2018). Artificial intelligence, automation, and work. In The economics of artificial intelligence: An agenda, pages 197–236. University of Chicago Press.

Aghion, P., Algan, Y., and Cahuc, P. (2011). Civil society and the state: The interplay between cooperation and minimum wage regulation. Journal of the European Economic Association, 9(1):3–42.

Ahrens, L. (2024). Labor market risks and welfare preferences: A bounded rationality approach. Socio-Economic Review, 22(1):327–348.

Alesina, A., Algan, Y., Cahuc, P., and Giuliano, P. (2015). Family values and the regulation of labor. Journal of the European Economic Association, 13(4):599–630.

Algan, Y. and Cahuc, P. (2009). Civic virtue and labor market institutions. American Economic Journal: Macroeconomics, 1(1):111–145.

Anelli, M., Colantone, I., and Stanig, P. (2021). Individual vulnerability to industrial robot adoption increases support for the radical right. Proceedings of the National Academy of Sciences, 118(47):e2111611118.

Angei, F., Brau, R., Balia, S., and Sulis, G. (2025). Perceptions on AI in the Italian labor market: Evidence from a survey experiment. AEA RCT Registry. RCT ID: AEARCTR-0016792. First published September 19, 2025.

Autor, D. (2024). Applying AI to Rebuild Middle Class Jobs. NBER Working Paper 32140, National Bureau of Economic Research.

Backus, P. G. and Esteller-Moré, A. (2017). Risk aversion and inequity aversion in demand for unemployment benefits. International Tax and Public Finance, 24(2):198–220.

Bick, A., Blandin, A., Deming, D. J., Fuchs-Schündeln, N., and Jessen, J. (2026a). Mind the Gap: AI Adoption in Europe and the U.S. NBER Working Paper 34995, National Bureau of Economic Research.

Bick, A., Blandin, A., Deming, D. J., and Schumacher, T. R. (2026b). What work does generative ai do? NBER Working Paper 35677, National Bureau of Economic Research.

Botero, J. C., Djankov, S., Porta, R. L., Lopez-de Silanes, F., and Shleifer, A. (2004). The regulation of labor. The Quarterly Journal of Economics, 119(4):1339–1382.

Brodeur, A., Cook, N., and Valenta, D. (2025). Perceptions of artificial intelligence and environmental sustainability. IZA Discussion Paper 18263, IZA Institute of Labor Economics.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative AI at Work. Quarterly Journal of Economics, 140(2):889–942.

Busemeyer, M. R., Gandenberger, M., Knotz, C., and Tober, T. (2023). Preferred policy responses to technological change: Survey evidence from OECD countries. Socio-Economic Review, 21(1):593–615.

Cardullo, G., Conti, M., Ricci, A., Scicchitano, S., and Sulis, G. (2024). On the emergence of cooperative industrial and labour relations. British Journal of Industrial Relations, 62(3):568–590.

Cattaneo, M. A., Gschwendt, C., and Wolter, S. C. (2025). How Scary Is the Risk of Automation? Evidence from a Large-Scale Survey Experiment. Journal of Economic Behavior & Organization, 235:107034.

Chen, W. X., Srinivasan, S., and Zakerinia, S. (2025). Displacement or complementarity?: The labor market impact of generative ai.

Cotofan, M. and Matakos, K. (2025). Recurring labour market shocks and stated and revealed preferences for redistribution. Journal of Economic Behavior & Organization, 238:107239.

Dalla Zuanna, A., Dottori, D., Gentili, E., and Lattanzio, S. (2024). An assessment of occupational exposure to artificial intelligence in italy. Questioni di Economia e Finanza (Occasional Papers) 878, Banca d'Italia.

Dechezleprêtre, A., Fabre, A., Kruse, T., Planterose, B., Sanchez Chico, A., and Stantcheva, S. (2025). Fighting climate change: International attitudes toward climate policies. American Economic Review, 115(4):1258–1300.

Di Tella, R. and Rodrik, D. (2020). Labour market shocks and the demand for trade protection: Evidence from online surveys. Economic Journal, 130(628):1008–1030.

Duval, R., Ji, Y., Papageorgiou, C., Shibata, I., and Spilimbergo, A. (2024). Preferences for labour regulation: endowments versus beliefs. Economic Policy, 39(119):549–606.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384(6702):1306–1308.

Engelhardt, C. and Wagener, A. (2018). What Do Germans Think and Know about Income Inequality? A Survey Experiment. Socio-Economic Review, 16(4):743–767.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal, 42(12):2195–2217.

Ferri, V., Porcelli, R., and Fenoaltea, E. M. (2024). Lavoro e intelligenza artificiale in italia: tra opportunità e rischio di sostituzione. Working Paper 125, INAPP.

Gallego, A., Kuo, A., Manzano, D., and Fernández-Albertos, J. (2022). Technological risk and policy preferences. Comparative Political Studies, 55(1):60–92.

Ge, S. and Zhou, Y. (2020). Robots, computers, and the gender wage gap. Journal of Economic Behavior & Organization, 178:194–222.

Golin, M. and Rauh, C. (2022). The impact of fear of automation. Cambridge Working Papers in Economics No. 2269, Faculty of Economics, University of Cambridge.

Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing information provision experiments. Journal of Economic Literature, 61(1):3–40.

Haslberger, M., Gingrich, J., and Bhatia, J. (2026). Rage against the machine? generative ai exposure, subjective risk, and policy preferences. Journal of European Public Policy, 33(11):2885–2912.

Humlum, A. and Vestergaard, E. (2026). Still waters, rapid currents: Early labor market transformation under generative ai. NBER Working Paper No. 33777, National Bureau of Economic Research.

Jaumotte, F., Kim, J., Koll, D., Li, E. Z., Li, L., Melina, G., Song, A., and Tavares, M. M. (2026). Bridging skill gaps for the future: New jobs creation in the AI age. IMF Staff Discussion Note SDN/2026/001, International Monetary Fund.

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.

Kuziemko, I., Norton, M. I., Saez, E., and Stantcheva, S. (2015). How elastic are preferences for redistribution? evidence from randomized survey experiments. American Economic Review, 105(4):1478–1508.

Lane, M., Williams, M., and Broecke, S. (2023). The impact of ai on the workplace: Main findings from the oecd ai surveys of employers and workers. OECD Social, Employment and Migration Working Papers, OECD Publishing.

Lergetporer, P., Schwerdt, G., Werner, K., West, M. R., and Woessmann, L. (2018). How information affects support for education spending: Evidence from survey experiments in germany and the united states. Journal of Public Economics, 167:138–157.

Margalit, Y. (2013). Explaining social policy preferences: Evidence from the great recession. American Political Science Review, 107(1):80–103.

Naumann, E. (2018). Tax constraints, social policy preferences, and support for redistribution. Societies, 8(3):46.

OECD (2026). Making AI Work: Why Investing in Skills Matters. OECD blog, 2 February 2026.

Okubo, T. and Wagner, A. F. (2026). Artificial intelligence and the demand for redistribution: Panel evidence. Available at SSRN.

Pew Research Center (2025). US Workers Are More Worried Than Hopeful About Future AI Use in the Workplace.

Shen, M., Xiao, Q., Chen, X., Fang, S., and Ngok, K. (2026). Generative AI, perceived job displacement, and policy preferences: Experimental evidence from china. Journal of Economic Behavior & Organization, 246:107535.

Stanford Institute for Human-Centered Artificial Intelligence (2026). The 2026 AI Index Report. Technical report, Stanford University.

Stantcheva, S. (2021). Understanding tax policy: How do people reason? Quarterly Journal of Economics, 136(4):2309–2369.

Thewissen, S. and Rueda, D. (2019). Automation and the welfare state: Technological change as a determinant of redistribution preferences. Comparative Political Studies, 52(2):171–208.