

Journal of Development Economics

AI Exposure and Policy Preferences: Experimental Evidence from Latin America --Manuscript Draft--

Manuscript Number:	DEVEC-D-25-00654R1
Article Type:	Registered Report Stage 1: Proposal
Section/Category:	Experimental Papers, credit, insurance
Keywords:	Technological adoption; Artificial intelligence; Policy preferences
Corresponding Author:	Yiming Cao, Ph.D The University of Hong Kong HONG KONG
First Author:	Yiming Cao, PhD
Order of Authors:	Yiming Cao, PhD Tian Chen, PhD Candidate Yican Liu, PhD Candidate Guillermo Rosas, PhD
Abstract:	Technological change can transform economies, generating both opportunities and disruptions that carry important political and policy consequences. Yet little is known about how people respond politically to new technologies at an early stage, when their eventual impacts remain uncertain. We study this question in the context of generative artificial intelligence (GenAI) through randomized online experiments in five Latin American countries. Participants are asked to perform a series of office-related tasks with or without access to a GenAI-powered assistant and then complete a survey on their beliefs about GenAI's potential influence and their preferences across a range of policy areas. To identify how these policy preferences are shaped by the beliefs individuals form, we implement an additional information treatment that experimentally manipulates perceptions of GenAI's potential impacts. The project sheds light on how people respond to novel technologies whose effects are uncertain and for which they have no prior experience.
Response to Reviewers:	

AI Exposure and Policy Preferences: Experimental Evidence from Latin America

Author statement

The following sections provide essential disclosures and acknowledgments pertaining to this manuscript. Details regarding the funding sources that supported this research, the ethical approval obtained for the study, acknowledgments of contributions from individuals or institutions, and a declaration of any artificial intelligence tools used during the preparation of this paper are presented below in accordance with the journal's submission guidelines.

Funding source: This experiment is supported by funding from the Carnegie Corporation of New York (ID: G-PS-24-61870). We also appreciate the financial support provided by the Weidenbaum Center on the Economy, Government, and Public Policy at Washington University in St. Louis; the Economic Research Grant at the University of Hong Kong; and the OpenAI Researcher Access Program (ID: 0000005377).

Ethical Approval: The experiment is approved by the IRB of the Human Research Protection Office at Washington University in St. Louis (ID: 202308187). And this experiment is registered at the AEA RCT registry (ID: AEARCTR-0014211).

Acknowledgement: We appreciate Shaokai Ding, Zhiyuan Li, Ronghua Shi and Tonglin Zhang for their excellent research assistance. We thank Timm Betz, Junyou Chen, Ted Enamorado, Guojun He, Ruilin Lai, Yu-Hsiang Lei, Fangzhou Lu, Henry Sauermann, Mingzhu Tai, and Anna Wilke, as well as seminar participants at the HKU Finance PhD Brownbag, HKU Economic Department, the Experimentation Webinar of Innovation Growth Lab, and the Political Data Science Lab and International Relations Workshop at Washington University in St Louis for their encouragement and helpful comments.

AI Assistance: We acknowledge the assistance of GPT-4o (OpenAI) and Copilot (Microsoft) in improving the readability of this paper. All errors are our own.

Journal of Development Economics
Registered Report Stage 1: Proposal

AI Exposure and Policy Preferences:
Experimental Evidence from Latin America

October 17, 2025

Abstract

Technological change can transform economies, generating both opportunities and disruptions that carry important political and policy consequences. Yet little is known about how people respond politically to new technologies at an early stage, when their eventual impacts remain uncertain. We study this question in the context of generative artificial intelligence (GenAI) through randomized online experiments in five Latin American countries with a total of 12,000 participants. Participants are asked to perform a series of office-related tasks with or without access to a GenAI-powered assistant and then complete a survey on their beliefs about GenAI's potential influence and their preferences across a range of policy areas. To identify how these policy preferences are shaped by the beliefs individuals form, we implement an additional information treatment that experimentally manipulates perceptions of GenAI's potential impacts. Overall, the project seeks to shed light on how people respond to novel technologies whose effects are uncertain and for which they have no prior experience—particularly in developing countries that may be more vulnerable to such disruptions.

Keywords: Technological adoption; Artificial intelligence; Policy preferences

JEL codes: O33; O38; O54; P00

Study pre-registration: AEA RCTR-0014211

1 Introduction

Recent developments in Generative AI (GenAI) have sparked widespread discussion about its potential to transform the economy and society. Existing studies have provided early evidence on how GenAI may influence productivity and reshape job opportunities across a broad set of tasks, occupations, and worker groups (Acemoglu, 2024; Felten, Raj and Seamans, 2021). These possibilities carry important implications for policy, yet we know little about what kind of policy citizens are likely to demand in response. This matters because the appropriate policy response depends not only on what may be economically optimal given the technology’s actual impact — which itself remains uncertain — but also on how individuals perceive that impact and the demands they express.¹ More broadly, this raises the question of how people respond when confronted with novel technological shocks whose consequences are unknown and for which they have no prior experience.²

In this study, we examine how exposure to GenAI shapes individuals’ perceptions of its capabilities — both with respect to potential productivity effects and possible implications for labor market outcomes — and how these perceptions translate into the policies they demand in response. We focus specifically on five Latin American countries — Argentina, Brazil, Chile, Colombia, and Mexico. As in many other regions, these countries are largely passive recipients rather than active developers of frontier AI technologies, and they often lack the institutional resources and state capacity needed to cushion the shocks that such technologies may generate.³ Examining how GenAI shapes political preferences in this context is therefore particularly relevant, yet systematic evidence on its effects in such settings remains scarce.

We conduct large-scale online experiments with 12,000 participants. In the main experiment, we randomize exposure to GenAI tools through simulated office tasks and

¹Demands may diverge from the economically optimal policy because individuals may not fully recognize the potential impact of the technology, or, even if they do, they may attribute labor market outcomes to other forces such as trade or immigration (Card, 2009; Rodrik, 2021; Wu, 2022).

²This is a fundamental issue because early reactions can shape policy responses that, in turn, influence technological trajectories. Yet such moments are rarely studied, largely because suitable settings are difficult to find: most evidence on the link between technology and politics comes from retrospective studies of shocks whose outcomes were already realized, and systematic data on how citizens respond under genuine uncertainty are seldom available (e.g. Caprettini and Voth, 2020; Molinder, Karlsson and Enflo, 2021).

³There are multiple reasons why Latin America is unprepared to deal with the effects of AI, including weak institutions (Acemoglu and Robinson, 2012), fragmented labor markets (Heckman and Pagés, 2000; Gong, Van Soest and Villagomez, 2004), ideological polarization (Campello and Urdinez, 2021), and limited redistributive capacity (Holland, 2018). Less than one-third (28%) of Latin Americans believe their countries are prepared to deal with AI (IPSOS, 2024), and governments in the region continue to lag behind other parts of the world in AI readiness (Nettel et al., 2024).

then measure subsequent changes in respondents' perceptions and policy preferences.⁴ All participants are asked to complete a set of office-related tasks, including text comprehension, numerical calculation, and business writing. In the treatment group, individuals are given access to an AI assistant powered by the most advanced large language model (LLM) available at the time of the experiment. In the control group, by contrast, individuals receive tutorial-style hints on how to perform the tasks, along with access to standard tools such as a calculator. To ensure comparability, both groups interact with their resources through an identical chatbox interface embedded in the task page — the only distinction being whether the assistance is AI-powered or not.

We use a follow-up online survey to measure participants' beliefs about GenAI's influence — both its potential productivity benefits and the risks it poses for job displacement — as well as the policies they demand in response. We elicit preferences over four broad bundles of policies: (i) those aimed at promoting or restricting the development and adoption of AI technologies; (ii) those that help workers adapt to the AI era, such as on-the-job training; (iii) those that mitigate AI-related risks through redistribution, such as universal basic income (UBI); and (iv) those that address other sources of labor market disruption, including trade and immigration. We elicit revealed preferences over these policies through incentivized, real-stakes questions following Stantcheva (2023) and Haaland, Roth and Wohlfart (2023). We will also conduct an additional survey seven days after the original intervention to assess whether the treatment produces effects beyond those that are immediately recorded.

In addition to estimating the average treatment effects of GenAI exposure on perceptions and policy preferences, we examine heterogeneity in these effects by considering the extent to which individuals' job tasks are likely to be affected by AI, as estimated in prior work (Felten, Raj and Seamans, 2019, 2021). We are especially interested in whether individuals in more exposed occupations will be more likely to perceive both the risks and benefits of the technology, and thus will exhibit larger shifts in their policy preferences.

To more precisely identify how perceptions of GenAI shape policy preferences, we conduct an additional experiment that manipulates these perceptions after participants randomly interact with AI. Specifically, participants are provided with information emphasizing either the potential productivity benefits of GenAI or the risks of job

⁴Baseline exposure to AI remains considerably lower in the countries we study than in advanced economies. However, we recognize that many participants are likely to have already encountered AI through features embedded in products or through news and public discussions. Our experiment is not designed to provide first-time exposure to GenAI, but rather to create randomized variation in the nature and extent of exposure. To this end, we carefully design tasks that are more complex and engaging than typical daily interactions with AI. Given lower baseline exposure and the relative complexity of our tasks, we expect that our treatment will generate meaningful variation in participants' interaction with and understanding of AI capabilities.

displacement.⁵ This treatment is implemented among a sub-sample of 6,000 participants after they complete the set of tasks with or without AI assistance, and is designed to alter the salience of the two perceptions that we hypothesize are formed through interaction with AI. We hypothesize that individuals primed with information about productivity gains will be more likely to support policies that promote AI adoption and help workers adapt, while those primed with information about displacement risks will be more likely to support restricting AI development and redistributive policies. This additional experiment allows us to identify how the specific beliefs about AI that individuals form shape their policy preferences following AI interaction.⁶

This paper contributes to several key research strands by examining the impact of GenAI on the attitudes and policy preferences of individuals. Our research offers a unique perspective on how individuals arrive at policy preferences as they interact with this emergent, revolutionary technology. First, we contribute to the growing literature on the labor market impacts of GenAI. Scholars are beginning to explore the economic implications of this technology, viewing it as a new wave of automation that can both complement and substitute for labor (Acemoglu, Ozdaglar and Siderius, 2024). Early evidence suggests that GenAI may affect occupations requiring higher cognitive skills that were previously thought to be less vulnerable to automation (Felten, Raj and Seamans, 2019, 2021; Tomlinson et al., 2025). This may exacerbate workplace inequality and contribute to job polarization (Hui, Reshef and Zhou, 2024; Karunakaran et al., 2025). However, other studies highlight its potential to generate significant productivity gains, particularly for lower-skilled workers (Noy and Zhang, 2023; Brynjolfsson, Li and Raymond, 2025). Our research expands this body of work by exploring how individuals form beliefs about the direction of GenAI's impact after direct interaction and, crucially, how these perceived labor shocks translate into shifts in policy attitudes. Our findings have important policy implications for managing technological disruptions, as technology-induced economic shocks can heighten the risk of political instability and social unrest (Caprettini and Voth, 2020; Gallego and Kurer, 2022; Guriev and Papaioannou, 2022).

Second, we contribute to the literature on the political impact of technological adoption. Building on the canonical task-based framework of technological progress, which posits that advancements like automation and AI create labor market shocks by substituting for, complementing, and creating jobs (Acemoglu, 2002b; Autor, Levy and Murnane, 2003; Acemoglu and Restrepo, 2018, 2019, 2020), a growing body of research

⁵A placebo group receives irrelevant information about global entertainment developments.

⁶To assess the independent role of information, we include a group of 3,000 additional participants who only receive the information treatment without completing any tasks.

explores how these shocks influence political dynamics. In historical contexts, studies show that labor-saving technologies, such as the steam engine, can lead to technological backlash (Caprettini and Voth, 2020), while general-purpose technologies that expand labor demand can increase workers' bargaining power (Molinder, Karlsson and Enflo, 2021). More recent research links automation to shifts in political behavior, including increased support for redistribution (Thewissen and Rueda, 2019), for populist and radical right parties (Anelli, Colantone and Stanig, 2019, 2021; Milner, 2021), as well as to changing attitudes toward technological progress (Gallego and Kurer, 2022) and toward the speed of adoption of new technologies (Gallego et al., 2022).

We extend this literature by studying the political effects of an *emerging* technology, specifically GenAI, rather than mature technologies. Our findings also provide insight into the determinants of technological adoption from the perspective of individuals and workers, complementing theories on technological change that typically focus on supply-side incentives for firms and scientists (Acemoglu, 2002a; Hémous and Olsen, 2021).

Finally, our research contributes to the broader literature on how economic shocks influence individual beliefs and attitudes. Recent work has shown that shocks from financial crises and globalization can fuel extreme ideologies and political polarization (Autor et al., 2020; Gyöngyösi and Verner, 2022). Our study also builds on the literature on expectation formation, which explores how people synthesize past experiences with new information to form beliefs about the future (Malmendier and Nagel, 2011; Jiang et al., 2024; Bordalo et al., 2023; Coffman, Collis and Kulkarni, 2024; Graeber, Roth and Zimmermann, 2024; Bordalo, Conlon, Gennaioli, Kwon and Shleifer, 2025). By examining how individuals form expectations about GenAI in response to direct interactive experiences, we provide an empirical extension of these theories in the context of a transformative economic shock (Bordalo, Burro, Coffman, Gennaioli and Shleifer, 2025). Our findings can inform policy to deal with the inevitable propagation of GenAI. Specifically, we propose that individuals' perceptions of GenAI's impact on their employment directly influence their policy preferences. Those who perceive GenAI as a threat are more likely to support regulation and favor stronger redistribution, while those who view it as a productivity enhancer are more inclined to support its adoption and express weaker redistributive preferences.

2 Background

2.1 Generative Artificial Intelligence (GenAI)

A strand of literature in labor economics suggests that technologies can act as substitutes for or complements to labor, depending on the specific task components involved (Acemoglu, 2002b; Autor, Levy and Murnane, 2003; Acemoglu and Restrepo, 2018, 2019, 2020). This insight also aligns with political economy literature that links the threat of job displacement from technological progress to individuals' political preferences (Guriev and Papaioannou, 2022; Gallego and Kurer, 2022). Within this framework, GenAI represents a novel wave of skill-biased technological change that may either substitute for or complement various occupations, depending on task and skill requirements (Acemoglu, 2021). Unlike earlier technologies — such as industrial robots, which primarily displaced routine tasks — recent GenAI models, including those in the GPT family, affect even non-routine, high-skill occupations. These models demonstrate superior performance in programming and enhance productivity in creative and analytical work (Bubeck et al., 2023; Eloundou et al., 2023; Brynjolfsson, Li and Raymond, 2025). Occupations now considered at high risk of replacement include roles such as teachers and social scientists, which were previously viewed as less susceptible to automation (Felten, Raj and Seamans, 2021, 2023; Tomlinson et al., 2025). Though GenAI demonstrates its potential to replace non-routine, cognitively demanding jobs, recent research also points out that GenAI can benefit workers, particularly low-skill workers, by enhancing their productivity. Noy and Zhang (2023) finds that individuals' productivity improves in a writing task when they are assigned to use ChatGPT in an online experiment. Furthermore, in a randomized experiment among customer support agents, Brynjolfsson, Li and Raymond (2025) indicates that GenAI led to increased productivity, measured by issues resolved per hour. A heterogeneous analysis shows that these productivity gains are primarily driven by low-skill and new workers, suggesting that GenAI can accelerate the movement of newer workers along the experience curve.

2.2 Latin American Countries

We study the effects of exposure to AI performance in Latin America, specifically in Argentina, Brazil, Chile, Colombia, and Mexico. Latin American economies are highly segmented, with large informal sectors (where laborers are not connected to state fiscal or welfare institutions) that coexist with well-organized formal sectors where state capacity is abundant and where many individuals are employed in firms that compete globally. In general, we expect the degree of exposure to AI to be extremely low in the informal

sector and much higher in the formal sector of the economy, but still lagging behind in exposure to AI compared to workers in more advanced economies. In particular, even workers in the formal sector of the economy are not routinely exposed to AI tools in their work (Di Battista et al., 2023; Nettel et al., 2024). For example, in a pilot survey we conducted in Mexico in January 2025, only about half of respondents reported having heard of ChatGPT. In short, while many Latin American formal-sector workers are not entirely unfamiliar with AI, relatively few have had extensive opportunities to use AI tools for work-related purposes prior to our study.

In addition, estimates of the degree to which different tasks can be performed by GenAI suggest that relatively high-skilled, mid-wage workers in the formal sector in Mexico are most vulnerable to job loss from AI replacement, whereas these same estimates suggest that AI would largely affect low-skill, low-wage workers in the US (Benítez-Rueda and Parrado, 2024). Gmyrek, Winkler and Garganta (2024) similarly suggests that highly-educated urban-based workers receiving relatively high incomes in the formal sector of Latin American economies — for example, in finance, insurance, or public administration — are much more likely to be exposed to AI in the immediate future. At the same time, persistent differences in digital access across formal-sector companies suggests that not all workers in the urban, high-skill, high-income category will be immediately exposed to the use of AI. Likely, only individuals in large multinational companies, core national public agencies, or prestigious universities have routine access and exposure to general-purpose AI like ChatGPT, and even these entities may lag behind in relatively-poorer countries in the region. Thus, Gmyrek, Winkler and Garganta (2024) estimates that both AI-driven job automation and productivity-enhancement potential in Latin America lags far behind the US, Canada, and Western Europe, and also trails China and India.

In short, Latin America provides an ideal setup to understand the effects of AI exposure on policy preferences. The region has lower exposure to AI than the US, Europe, China or India, which suggests ample availability of workers even in formal, globally-competitive sectors of the economy that have not incorporated AI to their daily routines. Moreover, workers within these sectors are potentially at higher risk of job loss than similarly-situated workers elsewhere. Finally, though welfare institutions in Latin America tend to support workers in the formal economy, to the detriment of informal workers, they have scant ability to insure against job loss or to provide retraining skills to unemployed individuals.

3 Theory and Hypotheses

The objective of this study is to examine how interaction with GenAI affects individuals' perceptions of its influence and how those perceptions shape their preferences for policies in response to the GenAI technological shock. We focus on two central dimensions of the perceptions that individuals might generate about GenAI, namely, its potential to generate productivity gains and the risks it creates of massive job displacement. Both dimensions have been widely discussed and empirically explored as potential outcomes of GenAI. For example, Noy and Zhang (2023) find that individuals assigned to use ChatGPT in a writing task demonstrate significantly higher productivity, while Brynjolfsson, Li and Raymond (2025) show that GenAI tools improve performance among customer support agents, especially low-skill and new workers. In contrast, Tomlinson et al. (2025) identify occupations, including teachers and social scientists, as increasingly susceptible to replacement by GenAI, highlighting growing concerns about job displacement. Based on these findings, we hypothesize that interacting with AI prompts individuals to form beliefs along these lines.

We study the effect of GenAI exposure on individuals' policy views. We examine three types of policy bundles people may demand upon the arrival of the AI shock:⁷

- a. Defensive policies aimed at restricting AI adoption and enhancing GenAI regulation. These reflect concerns about labor market disruptions, monopolization by private or foreign actors, and privacy risks. We assess how strongly individuals support government intervention to mitigate such harms.
- b. Active policies that help individuals adapt to GenAI development. Recognizing its potential for productivity gains, citizens may favor government-provided career training and the inclusion of GenAI education in schools. This bundle captures support for proactive government involvement in AI adoption.
- c. Passive policies that do not directly target GenAI but address its distributional consequences. For instance, individuals may support or oppose redistributive policies, such as universal basic income (UBI), in response to anticipated economic losses and gains from GenAI adoption.

In addition to these three bundles, we examine whether GenAI influences attitudes toward broader economic policies, such as immigration and globalization, which are

⁷We focus on policy preferences rather than political behaviors (e.g., protesting or voting). Such behaviors could be affected by AI exposure, but we do not believe that the kind of short exposure we will provide would trigger them. Policy preferences, however, are closely linked to voting behavior as suggested by literature on "policy voting" (Page and Brody, 1972; Fowler et al., 2020, 2024a,b).




Figure 1: Expected reactions to perceptions of job loss and productivity gains from interactions with GenAI

often erroneously blamed as the cause of job disruptions that are in fact driven by technological change (Card, 2009; Rodrik, 2021; Wu, 2022). Figure 1 summarizes how GenAI interaction affects the outcomes of interest, conditional on potential perceptions induced by such interaction.

3.1 Hypotheses

We propose a first set of hypotheses to investigate whether interaction with GenAI leads to stronger perceptions of replacement risks and productivity gains. In other words, we suggest that change in policy attitudes will follow from interaction with GenAI through two potentially distinct perceptual mechanisms: replacement risks vs productivity gains.

- **H1a:** Interaction with GenAI increases individuals' perceived likelihood of being replaced by GenAI.
- **H1b:** Interaction with GenAI increases individuals' perceived likelihood of experiencing productivity gains in their jobs.

As shown in Figure 1, our proposed mechanisms suggest that interaction with GenAI will lead to stronger support for GenAI regulation (defensive policies), stronger opposition to GenAI (active policies), and increased support for redistribution (passive policies) if individuals perceive GenAI as a threat to their jobs. Conversely, perceived

productivity gains from GenAI may produce the opposite effects. Because the impact of GenAI depends on the relative magnitude of the two perceptual mechanisms, we lack firm expectations about the direction of the average treatment effect of GenAI on policy preferences. This effect can be positive, negative, or nil depending on whether average perceptions of job replacement are weaker than, stronger than, or equal to perceptions of productivity gains.

The second set of hypotheses tests whether perceived risks of job replacement and/or perceived productivity gains are true mediators of the effect of interaction with GenAI on policy preferences. The following hypotheses examine the causal relationship between individuals' perceptions of GenAI and their policy preferences.

- **H2a:** Individuals for whom perceptions of job loss are more salient will increase support for defensive policies, decrease support for active policies, and increase support for redistributive policies after interacting with GenAI.
- **H2b:** Individuals for whom perceptions of productivity gains are more salient will decrease support for defensive policies, increase support for active policies, and decrease support for redistributive policies after interacting with GenAI.

4 Experimental design

4.1 Respondent sample

We plan to conduct an online randomized experiment with representative samples from five Latin American countries in collaboration with *NetQuest-Bilendi*.⁸ We conduct the experiment online for two main reasons. First, from a practical perspective, recruiting respondents in the field and tracking their beliefs and attitudes over time would be extremely challenging. Our sample providers have years of experience conducting surveys in Latin American countries and large panels of registered and vetted participants, which enables us to obtain participant samples that are broadly representative of the population connected to the internet. Additionally, we use quotas in the survey platform to ensure our samples closely match the national distribution of key characteristics, such as educational background and age. Second, online experiments help mitigate experimenter effects and social desirability bias, as suggested by recent research (Mummolo and Peterson, 2019; Haaland, Roth and Wohlfart, 2023). In any case, recent research in economics suggests that online-panel samples exhibit similar response patterns and

⁸<https://www.netquest.com>

characteristics — especially in terms of political views and knowledge — to those of offline probability-based samples (Coppock and McClellan, 2019; Grewenig et al., 2023; Haaland, Roth and Wohlfart, 2023).

Our sample includes 12,000 individuals from Argentina, Brazil, Chile, Colombia, and Mexico. We select these countries because they are the largest economies in the region. Additionally, these countries have highly internationalized formal labor markets and engage in frequent trade with two economic superpowers, the United States and China, that are at the forefront of development of AI. From this perspective, conducting our experiment in multiple countries enhances the external validity of our results and provides important policy implications for other developing countries in this region and beyond, as these countries are very likely to be affected by future international AI-driven economic shocks.

Specifically, our sample includes individuals of working age (18 to 65), and we set quotas consistent with the distribution of sex, age, and socio-economic status (education and income) among national populations with access to the internet. It is worth noting that most respondents will possess computer literacy, as participation in the online experiment requires access to a computer or mobile device. These individuals are more likely to be urban-based and working in the formal sector of the economy. We argued above that these are the individuals most likely to be affected by AI-driven economic changes, and we are aware that our sample has the potential to under-represent workers in the informal sector of the economy, which we acknowledge as a potential limitation.

It is worth noting that the sample size exceeds the rule-of-thumb suggested by Haaland, Roth and Wohlfart's (2023) power requirement for an information provision experiment. According to Haaland, Roth and Wohlfart (2023), to achieve at least 80% power to detect a treatment effect of 15% of a standard deviation of the outcome variable, each treatment arm requires a minimum of 700 respondents. In our experimental design, we adopt a conservative approach to sample size, selecting a larger sample than Haaland, Roth and Wohlfart's recommended benchmark. As illustrated in Figure 2 and summarized in Table 1, the experiment is organized into three distinct sub-groups. The first sub-group investigates the effects of completing tasks with GenAI, assigning 3,000 individuals to the AI-exposure treatment and another 3,000 to the control group, altogether 6,000 respondents. The second sub-group is used to examine our proposed mechanism. This sub-group consists of six treatment arms, each combining one of three information interventions with one of two task interventions, with 1,000 individuals allocated to each arm (6,000 total). Notably, every sub-group in our design exceeds the threshold of 700 participants when we pool together respondents from all countries. This substantial sample size is intended to enhance the statistical power of our analysis and

Table 1: Summary of the Experimental Procedure

Time	Randomization, intervention, or survey	% of Participants	Intervention or survey details
<i>Intervention period</i>			
Day 1	Initial survey	100%	Online Appendix
Day 2	Randomization & interventions		
	Randomization 1:		
	Randomize participants into three groups		
	Experiment T: Exposure treatment (doing tasks)	40% of all	
	Experiment TI: Joint treatments (info. & tasks)	40% of all	
	Experiment I: Information treatment only	20% of all	
	Randomization 2 (for T and TI):		
	Assign task intervention:		Appendix A.3
	(<i>Control</i>) Use of non-AI tools to finish tasks	50% of T&TI	
	(<i>Treatment</i>) Use of AI tools to finish tasks	50% of T&TI	
	Randomization 3 (for I and TI):		
	Assign information intervention:		Appendix A.1
	(<i>Control</i>) Irrelevant-information	30% of I&TI	
	(<i>Treatment</i>) How GenAI replaces jobs	35% of I&TI	
	(<i>Treatment</i>) How GenAI increases productivity	35% of I&TI	
Day 2	Final survey	100%	Online Appendix
<i>Post-intervention period</i>			
Day 8	Follow-up survey	100%	Online Appendix

support robust, reliable conclusions from the experiment.

4.2 Organization of the experiment

To evaluate our hypotheses, we structured the experiment into three sub-experiments. The first experiment tests how GenAI interaction shapes individuals' perceptions of GenAI and their policy preferences. In the second experiment, which follows the GenAI interaction intervention, we manipulate the salience of two perceptions of GenAI by randomly providing information on GenAI's potential for job replacement and its potential productivity gains. That is, the second intervention is an "information treatment" that seeks to introduce exogenous variation in the salience of perceptions because we understand that perceptions may be affected by additional factors that we

do not control; barring this additional manipulation, we would be at risk of inferring effects of perceptions that are in fact produced by some other attitude or by some other individual characteristics, including socio-economic factors. In the third experiment, we isolate information effects by conducting the information treatment only. The randomization and experimental process are summarized in Figure 2 and Table 1, with further details documented below.


Experiment T: Task treatment only We perform this experiment on 6,000 individuals to evaluate the impact of direct interaction with GenAI tools, and to test our main hypotheses (H1a and H1b). Participants are assigned a multi-step task and randomly divided into two conditions: the control group (50%, about 3,000 individuals), which completes the task without GenAI assistance, and the treatment group (50%, about 3,000 individuals), which completes the task with GenAI support. This design allows us to investigate the effect of GenAI exposure on individuals' perceptions of GenAI and their policy preferences.

Experiment TI: Sequential Assignment to Tasks and Information Frames We perform the second experiment on approximately 6,000 individuals to test the hypothesized mechanisms underlying AI's influence on policy preferences. Participants in this group are randomly assigned to receive information about GenAI's potential job displacement effects and productivity gains *after* completing the multi-step task. In a first step, the sample is evenly split between those who complete the task with GenAI support (3,000 individuals) and those that complete the task without such support (3,000). Following task completion, participants are randomly exposed to one of three types of information: 30% (1,800) receive irrelevant information (a news clip on the entertainment industry), 35% (2,100) receive information about GenAI's job replacement effects, and 35% (2,100) receive information about GenAI's productivity gains. By manipulating the salience of individuals' perceptions of GenAI through the information provision experiment in a sequential experimental design, we test the proposed mechanism implied by hypotheses H2a and H2b.

4.3 Interventions

To test our hypotheses, we conduct a randomized online experiment in which participants receive targeted instructions and engage in tasks designed to enhance their exposure to AI technologies. The experimental design is shown in Table 1 and Figure 2. First, we conduct a baseline survey on Day 1 to capture individuals' political

Figure 2: Experiment Design



and economic characteristics, as well as their prior exposure to, perceptions of, and beliefs about AI. On Day 2, we assign participants to complete the multi-step task and information treatments.

Exposure treatment: Using GenAI to Solve a Multi-Step Task Before assigning participants into treatment and control groups, we ask them to complete a simple task without any assistance, as a warm-up exercise. After this step, our main treatment starts by introducing half of the participants to an instructional video on how to effectively prompt ChatGPT. In the control group, individuals are asked to watch an instructional video that explains how to efficiently read emails and work documents, and how to

politely communicate with their boss via email.⁹ Following the video, individuals are asked to solve a complex multi-stage task. In the treatment group, individuals are encouraged to use a GenAI assistant powered by GPT-5 to solve the task. To minimize attrition, we have designed an AI interface within the Qualtrics survey platform that allows respondents to access ChatGPT without navigating to an external page.¹⁰ Respondents in the control group will be required to complete the same task. In order to promote diligent participation, we ask respondents in both groups to advance as much as they can in solving the task before hitting a time-limit of 12 minutes.

To ensure that respondents are exposed to the broad applications of GenAI, we design a complex task composed of multiple stages. In this task, respondents are asked to resolve a consumer claim as an employee at a chip company. We begin by presenting a business email from a consumer complaining about a delay in the company's product. The email is intentionally wordy and lacks clear logic. In the first sub-task, respondents must identify the main concern expressed in the customer's email. In the second and third sub-tasks, respondents are required to develop a plan for arranging processor production across different factories, which involves calculating production costs and determining the most efficient production strategy based on the provided data. In the final stage, respondents are informed that the company has chosen a plan and are asked to write a business email to the import licensing office of a foreign country using information from a provided guidance document. We expect that by completing this complex task, individuals in the treatment group will better understand how GenAI can be applied in various areas, including reading comprehension, mathematical calculations, and writing. A detailed description of these tasks appears in Appendix A.3. Appendix A.4 provides a detailed description of the interface and presents the AI-based interface used for the treatment group.

To account for potential differences in engagement between the control and treatment groups, we ensure that participants in the control group remain fully engaged and follow instructions carefully by providing them with a bespoke rule-based non-AI-powered chatbot with interface similar to the AI chatbot used by the treatment group. The rule-based chatbot includes a calculator and accessible tips to complete the task, but purposefully avoids interaction with GenAI. The interface is described in Appendix A.5, and Figure A2 displays a visual of this interface.

Although our treatment introduces a relatively small “dose” of interaction with GenAI, its primary purpose is to introduce random and measurable variation in indi-

⁹We will collaborate with a video artist to create an engaging video for our respondents. The content will focus on a general introduction to ChatGPT and prompt engineering.

¹⁰See Appendix A.4.

viduals' experiences with this revolutionary technology. This allows us to examine how individuals form or alter perceptions about the impact of GenAI.

Information treatment: AI-related information To assess the explanatory potential of the hypothesized mechanisms, we additionally include an information treatment in which we ask respondents to read different statements. Three groups receive different information frames: one group receives information about AI's effect on job replacement, a second group receives information about AI's effect on increasing productivity, and a third group receives a placebo statement about the entertainment industry.¹¹ The full statements are provided in Appendix A.1. To ensure participants read the statements carefully, we include as an attention check a multiple-choice question about content immediately after reading the information treatments.

4.4 Survey methodology

In this section, we illustrate the methodologies used in our survey. As shown in Table 1, we have three rounds of surveys in our experiment: an initial survey, a final survey, and a follow-up survey. Because we conduct our experiment *online* and *anonymously*, we do not expect large experimenter effects throughout our survey, and participants are expected to offer their actual attitudes and opinions throughout our experiment. However, we still provide a clear design to elicit people's actual attitudes towards AI and to reduce experimenter effects. In the survey questions, we strictly adhere to prior guidance on survey design (Haaland, Roth and Wohlfart, 2023; Stantcheva, 2023):

Initial Survey: Basic Information About Participants Prior to our experimental interventions, we gather basic socio-demographic information about the participants. We also inquire about their knowledge of, attitudes toward, and beliefs about AI, particularly the technology of GenAI. To mitigate potential experimenter effects, we include several *Decoy questions* that inquire about respondents' attitudes toward other technologies, such as biotechnology or financial technology.

Final Survey: Eliciting Outcomes of Interest Following experimental interventions, we collect a comprehensive description of participants' attitudes toward AI and AI-related policies. Specifically, the final survey consists of three main sections. First, we inquire about individuals' attitudes toward AI, focusing on their beliefs regarding how

¹¹We also conduct an ancillary experiment involving only the information treatment, which serves as an additional benchmark for Experiment TI. See Appendix A.6 for details.

AI will affect the general labor market and their own jobs. Next, we ask about attitudes toward government policies related to AI, such as whether respondents would welcome active governmental support to propagate AI through the country’s economy. The final section covers policy preferences indirectly related to AI attitudes, including views on redistribution, immigration and globalization. Like the initial survey, we also include a *decoy question* on policies related to renewable energy and online learning platforms in the final survey, when asking respondents about their opinions on AI policies.

Follow-up Survey: Robustness and Impact of GenAI Effects One Week After Treatment We conduct a follow-up survey one week after our intervention. We design the follow-up survey for two reasons. First, for the sake of understanding how robust our findings are, we can re-examine our main conclusions across multiple rounds. Second, we can test the impact of our interventions beyond their immediate effects. To mitigate experimenter effects, our follow-up survey also includes obfuscating questions about other technologies.

5 Data

5.1 Timeline

We conducted a pilot experiment to provide proof of concept and basic knowledge about the effectiveness of survey question wording in January 2025. We will field our main online survey experiment (three rounds) and collect the bulk of the survey data within a month of acceptance of our pre-registered design. Further details on the survey data we will collect appear in the Online Appendix.

- **January 2025 - February 2025:** Pilot experiment.
- **February 2025 - August 2025:** Finalization of the RCT design based on lessons from the pilot experiment.
- **Early November 2025:**¹² Recruitment of 12,000 participants from Argentina, Brazil, Chile, Colombia, and Mexico.
- **Mid November 2025:** Initial survey on the field. This is the *first* dataset we need for our analysis.

¹²We keep this date as a placeholder. We will be able to initiate this stage within weeks of acceptance of our pre-registered experiment.

- **Mid November 2025** Experimental interventions:
 - Main Intervention: Multi-stage task with (without) AI tools
 - Joint intervention (testing mechanism): Including both multi-stage task and information frames about GenAI.
 - Information intervention: Information frames about GenAI.

We plan to collect information about the task performance, attitudes and opinions of participants during the experimental intervention. This is the *second* dataset we need for our analysis.

- **Mid November 2025:** Final survey, including questions on political attitudes (redistributive preferences, political opinions, etc.), knowledge and usage of AI, etc. This is the *third* dataset we need for our analysis.
- **Late November 2025:** Follow-up survey after one week. This is the *fourth* dataset we need for our analysis.

5.2 Data Collection and Processing

We collect multiple datasets to support our analysis. To summarize, there are four major datasets generated by the participants, including the initial survey, the intervention (tasks), final survey, and follow-up survey. Table 2 displays the variables that we will collect throughout our experiment.

5.2.1 Measurement Concerns and Variable Construction

This section extends our discussion from Section 4.4, where we described how we design an appropriate survey questionnaire to elicit participants' *actual* opinions, attitudes, and knowledge related to GenAI. In this section, we explain, in detail, how we construct a dataset based on the variables listed in Table 2. Additionally, we provide an Online Appendix with additional documentation and the full text of all survey questions in our experiment.

Initial Survey: Participants' Background Information and Prior Attitudes and Knowledge About AI The variables we collect include: (i) basic personal information, including nationality, gender, marital status, religious affiliation, household demographics, education level; (ii) employment information, including occupation, employment status, employment industry, and income status; and (iii) political attitudes, including party

Table 2: Summary of the Variables Collected in Our Experiment

Variables	Dimensions
<i>[Prior survey]</i>	
Basic information	
Demographics (Gender, education, etc.)	-
Political affiliation & attitudes	-
Employment	-
(i) <i>Industry in which respondent works/has worked</i>	
(ii) <i>Occupations and AI-exposure score calculated</i>	
(Prior) Knowledge about AI	Knowledge
(i) <i>AI definition</i>	
(ii) <i>AI products & firms</i>	
(iii) <i>AI algorithms & technological basis</i>	
(Prior) Belief about AI's impact on macroeconomics	Beliefs
(i) <i>labor market</i>	
(ii) <i>macroeconomic growth</i>	
(Prior) Belief about AI's impact on jobs	Beliefs
(i) <i>capabilities of AI in different dimensions</i>	
(ii) <i>job displacement by AI in the economy</i>	
<i>[During intervention]</i>	
Times of interaction with AI tools (Treat only)	Behaviors
Time (in minutes) of finishing the tasks	Behaviors
Quality of submitted answers	Behaviors
<i>[Final survey]</i>	
Perceptions about AI's impact on the economy	Beliefs
Perceptions about AI's quality, capabilities, & power	Beliefs
(i) <i>AI's quality and capabilities</i>	
(ii) <i>AI's drawbacks and limitation</i>	
Main variable of interests: AI-related policy preferences	
Defensive policies: Attitudes towards AI-regulation	Attitudes & Opinions
(i) <i>general restrictions in AI usage</i>	
(ii) <i>restriction about AI-adoption in different public sectors</i>	
(iii) <i>AI-related tax preferences</i>	
(iv) <i>AI-related innovation supported by government</i>	
(v) <i>data-intensive innovation and data privacy regulation</i>	
Active policies: Attitudes towards AI-adaptation	Attitudes & Opinions
(i) <i>general attitudes towards AI-adaptation in their jobs</i>	
(ii) <i>attitudes towards on-the-job training about AI</i>	
(iii) <i>attitudes towards education for AI adaptation in schools</i>	
Passive policies: Redistribution and globalization	Attitudes & Opinions

Continued

- Redistributive preferences
 - (i) *taxing high-income population*
 - (ii) *government programs & social insurance*
- Attitudes towards globalization & immigrants
 - (i) *international cooperation with AI advanced countries*
 - (ii) *foreign direct investment*

[Follow-up survey]

Knowledge about AI-related technology	Knowledge
(i) <i>AI definition</i>	
(ii) <i>AI products & firms</i>	
(iii) <i>AI algorithms & technological basis</i>	
(Reported) daily usage of AI	Behaviors
Attitudes towards AI-related policies	Attitudes & Opinions
(i) <i>AI adoption supported by government</i>	
Decisions in political participation	Behaviors
Interests in learning AI-related knowledge	Behaviors

affiliation and political ideology (left-right self-placement). Please see questions from Category A in the initial survey appendix (Online Appendix A.7).

We additionally elicit respondents' prior knowledge, beliefs, and attitudes towards AI. As shown under Category C in Online Appendix A.7, we collect information about (i) knowledge of AI, including basic technology related to AI, AI-related products, etc.; and (ii) beliefs about the economic impact of AI, as well as AI applications in jobs.

Experimental Intervention During our experiment, we collect and measure a number of behaviors. In particular, we register how many times participants interact with either the AI or task-manager chatbox, time employed in finishing the tasks, and the quality of submitted answers.

Final survey: Perceptions about AI The first primary outcome variable we examine is individuals' perception and understanding of AI. Our survey captures this through three key dimensions: (i) beliefs about how AI influences the labor market and their jobs, which are our proposed mechanisms; and (ii) general perceptions of AI's quality, capabilities, and potential. Wording for these items appears in Sections A through B of Online Appendix A.9.

Final Survey: Policy Preferences The second set of primary outcomes comprises individuals' policy preferences. In our experiment, we focus on how AI adoption

alters participants' attitudes towards public policies that ameliorate the risks imposed by AI or multiply its potential benefits. As we discussed in our theoretical framework and hypothesis, our treatment could lead to change in participants' perceptions about AI-driven technological change, thus altering their perceived economic risks. To test our theory, we classify policy preferences into three categories: AI regulation (*defensive policies*), adaptation to AI (*active policies*), redistribution (*passive policies*), and broader economic policies related to immigration and globalization. The first two dimensions are directly correlated with people's perceptions about AI, and the latter two capture how people understand the potential economic risks/changes associated with AI-adoption. We now discuss these four categories (Category C and D in Online Appendix A.9 contain exact wording for these survey items.).

First, to evaluate whether our treatment influences participants' perceptions of AI, particularly its potential negative impact, we first elicit attitudes toward policies that aim to regulate AI, what we call "defensive policies." Specifically, we assess individuals' willingness or reluctance to adopt AI technologies across various public sectors, including education, public administration, healthcare, and public safety. For instance, question D.6 directly addresses participants' views on AI implementation in public safety. We also measure participants' attitudes towards taxation on AI. Additionally, we measure individuals' willingness to embrace AI innovation, another example of defensive policy. This category of policy preferences includes two key dimensions: (i) whether the government should actively support AI-related innovation, and (ii) related considerations such as the government's role in ensuring data protection. For example, in question C.2, we ask whether the government should spend more funding on AI-related innovation. In C.6, we ask whether the government should allow companies to use private data for AI training.

Second, we further examine public perceptions about adaptation to AI, focusing on a set of initiatives referred to as "active policies." Specifically, we explore participants' attitudes toward government support of on-the-job AI training and toward integration of AI-related education into school curricula. For example, in question C.7, we ask the participants whether their firms should provide training to help them adapt to AI systems.

Third, we assess attitudes toward a set of "passive policies" that, although not directly related to AI adoption or innovation, may mitigate its potential effects. In our theoretical framework, we hypothesize that individuals perceive economic risks arising from technological change driven by AI adoption. For instance, views on income redistribution or on the desirability of a Universal Basic Income (UBI) scheme may shift depending on participants' expectations about AI's potential impact on job

security. We include a series of questions regarding people's redistributive preferences and their attitudes toward UBI, such as questions D.2 and D.5 in the final survey. To complement these attitudinal measures, we include a real-stakes question (C.14) that invites participants to enter a lottery and allocate potential winnings among real-world organizations supporting different policy approaches or to themselves.

In addition to defensive, active, and passive policies, we gauge participants' attitudes toward openness to international cooperation, trade and globalization, and immigration. As we mentioned above, people often attribute technology-induced economic disruptions to factors such as migration and globalization, so one might expect that exposure to AI could lead to changes in attitudes toward immigration and trade.

To facilitate our analysis, we code all variables in the same direction: items that measure perception about AI and attitudes towards AI-related policies are coded so that large values correspond to greater *acceptance* of AI, stronger beliefs in the *positive* impact of AI on the economy, and greater *support* for policies about AI-adoption or AI-innovation.

Follow-up survey In our follow-up survey (Online Appendix A.10), which we have scheduled a week after the final survey, we capture participants' knowledge about AI, as well as perceptions and attitudes towards AI and AI-related policies. We also design behavioral questions to elicit participants' potential interest in studying AI-related technology.

5.3 Potential deviations from the intended sample

Attrition We anticipate that the sample in our baseline survey will differ from that of our final and follow-up surveys as a result of two sources of attrition:

- (i) Some participants will not finish the entire intervention phase or will not take the final survey after the experiment.
- (ii) Some participants will not take the survey seriously.¹³

Were this to occur, we plan to exclude these observations from our analysis and compare the baseline characteristics of experimental subjects who dropped out with those who did not (both within and across treatment and control groups). We will conduct balance tests of socio-demographic variables by looking at difference of means across participants in the treatment and control groups, using respondents in both the baseline and final surveys. Although we do not anticipate any differences in the distribution of

¹³We will identify possibly-inattentive respondents in our survey through attention checks.

characteristics between the baseline and final survey samples, we follow the strategies outlined in Lee (2009) to address attrition bias concerns if there are any unbalanced characteristics between the treatment and control groups.¹⁴

Non-compliance Non-compliance may be a concern, particularly for experimental participants assigned to use AI tools to complete tasks, as some participants in the treatment group may choose not to utilize the AI tools, while some participants in the control group may use external AI. The presence of never-takers or always-takers in our treatment group would only reduce the magnitude and statistical significance of the estimated effects.¹⁵ Therefore, we do not expect non-compliance to lead to an overestimation of GenAI's potential effect on policy preferences.

5.4 Pilot Study

In January 2025, we conducted an online pilot study with 1,235 participants in Mexico to examine how AI information and AI exposure (actual usage) affected policy preferences. The purpose of this pilot was to identify technical issues within our survey and to enhance our understanding of the sample population. We do not report the results of this pilot in the paper because we used a different experimental design,¹⁶ making it difficult to draw conclusions useful for the experimental design used in this paper. Nevertheless, we obtained important insights that will be helpful for the formal experiment.

First, the pilot results revealed that individuals in Latin America are comparatively less exposed to AI, especially ChatGPT. In the baseline survey for the pilot study, we asked respondents whether they had heard about various online products. Among the 1,235 total participants, only 48% (594) had heard of ChatGPT. In comparison, 1,184 individuals (96%) had heard of Netflix, 1,177 (95%) had heard of YouTube, and 1,179 (95%) had heard of Facebook. It is important to note that these figures only indicate awareness

¹⁴In our pilot study, we did not observe a huge attrition rate *in the short term*. Only about 15% participants dropped out after a two-day intervention in our pilot study. We further discuss the issue of attrition in the pilot study in the next section.

¹⁵We also took steps to minimize the likelihood of both types of non-compliance. To reduce the incidence of non-compliance in the treatment group, we informed all respondents that they would be entered into a lottery based on their accuracy and performance during the tasks, thereby providing participants in the treatment group with additional incentives to use AI, and participants in the control group to use Task Hints. To deter non-compliance in the control group, we implemented a warning system: if participants attempt to leave the current page (to use external AI), a warning message appears.

¹⁶In the pilot, we first provided participants with general AI information, then showed them individual scores that captured the degree to which AI could do tasks typical in their jobs, and finally administered the AI exposure treatment. We found strong negative effects of the information treatment but not of AI exposure. Moreover, the negative information treatment reduced participants' willingness to use ChatGPT (measured by time spent using AI) by around 50%.

of the product, not actual usage. Furthermore, according to data from the 2017, 2020, and 2023 editions of Latinobarómetro, overall participation in the digital economy in Latin America remains relatively low. An average of 67.8% respondents reported that they had never made an e-commerce purchase, with only a few countries — such as Colombia and Chile — showing comparatively higher usage (Latinobarómetro Corporation, 2025). From this perspective, we are less concerned about ceiling effects that could prevent us from detecting significant treatment effects.

Second, the pilot results confirm that both the willingness to use AI and the actual usage of AI are high among individuals in the AI exposure group. We found that among the 568 individuals provided with AI tools,¹⁷ only 41 (7%) did not use ChatGPT even when given encouragement, and 326 (57%) used the AI tools more than once. Understanding the actual number of ChatGPT users is important for our experiment, as we employ an encouragement design, meaning that we will estimate an intention-to-treat (ITT) effect rather than an average treatment effect (ATE). We attribute this high level of usage to the careful design of our in-survey GPT interface, which clearly piqued respondents' interest. We also verified whether respondents used the tool to solve the task or merely for casual chat; fortunately, we found that almost all of them used it exclusively for task-related purposes.

Third, we found that the attrition rate is not high between the first and second waves of the pilot survey, even though the tasks assigned to respondents were quite time-consuming. In the pilot, participants first completed a baseline survey and a business email writing task, and then, three days later, they received a second survey that required them to respond to a consumer request. On average, each task took around 10 minutes to complete. Despite the time commitment, only about 15% of individuals dropped out after the first wave. This attrition rate likely represents a worst-case scenario for our formal experiment, as the baseline survey in the formal experiment will not include any task requirements.

Most importantly, we refined many issues found in the pilot study in the new experimental design presented in this paper. First, we designed a more challenging task, as described in Section A.3, because most respondents found our previous tasks too easy. We believe that increasing task difficulty can help participants better understand the capabilities of AI in the exposure experiment. Instead of executing several small, independent tasks as in the pilot study, the formal experiment involves a single, more complex task within one survey. This approach increases the likelihood that respondents will engage with the AI and reduces the attrition rate associated with multiple survey waves. In addition, we refined our survey questions by adding more “obfuscated items”

¹⁷This number excludes individuals who did not pass the attention check in the baseline survey.

and a real-stakes question (C.14) to mitigate potential experimenter demand effects. Additionally, we revised the questions to ensure they are easier for respondents to understand, based on guidance from Stantcheva (2023).

6 Analysis

We first test the impact of our three different interventions by estimating a series of linear regression models on the outcome variables collected in our final survey. First, we estimate the direct impact of interacting with AI. As we hypothesize that interacting with AI may shape people's beliefs about AI's capability of job replacement or productivity increase, we run the following regression:

$$P_i = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{AI Task}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_i \quad (1)$$

in which P_i stands in for two outcome variables indicating respondents' perceptions about the effects of AI (i.e., as a productivity enhancer or as a job killer), and $1_{\text{AI Task}_i}$ is an indicator variable that equals one if individual i is in the group that has access to GenAI tools. X_i is a set of control variables from our baseline survey that includes gender, age, left-right ideological self-placement, employment status, family income level, and education level (see Section 5.2).

As we aim to test how the use of GenAI tools influences respondents' policy preferences, our initial empirical analysis also calls for estimation of a second regression, as follows:

$$Y_i = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{AI Task}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_i \quad (2)$$

Here, Y_i stands in for each of a number of outcome variables that capture participant i 's preferences towards AI-related defensive, active, and passive policies. Regression model (2) will be based on the subsample of respondents in the AI-treatment arm (Experiment T in Section 4.2).

As outlined in the theory and hypotheses section, the direction of treatment effects may vary depending on how individuals form perceptions about the effects of AI. Consequently, a simple regression model with a single treatment dummy is insufficient to uncover the underlying mechanisms or support the broader conclusions of our study. It merely captures the average treatment effect across a representative sample, without accounting for heterogeneity in responses or the pathways through which these effects operate. Thus, we provide additional analysis on mechanisms based on the information

provision experiment.

Based on the subsample of respondents subjected to Experiment TI (Section 4.2), we test whether there are differences in policy preferences among those randomly exposed to different information frames and different task accomplishment tools. To do so, we estimate the following regression:

$$Y_i = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{Info. Prod. Gain}_i} + \tilde{\beta}_2 1_{\text{Info. Job Loss}_i} + \tilde{\beta}_3 1_{\text{AI Task}_i} + \tilde{\beta}_4 1_{\text{Info. Prod. Gain}_i} \times 1_{\text{AI Task}_i} + \tilde{\beta}_5 1_{\text{Info. Job Loss}_i} \times 1_{\text{AI Task}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_i \quad (3)$$

Based on our empirical analysis, we test our main hypotheses related to participants' perceptions about AI and attitudes towards AI-related policies based on data collected in our final survey which, as mentioned in Section 5.2.1, includes variables related to perceptions about AI and attitudes towards AI-relevant policies that are coded so that higher values correspond to more support for defensive policies, less support for active policies, and more support for passive policies.¹⁸ Thus, we interpret our empirical results as follows:

Expectations about perceived job loss risks or productivity gains (H1a and H1b) An estimate of $\beta_1 > 0$ in regression model (1) would confirm participants' concerns about the possibility of job replacement and of productivity gains when P is, respectively, an outcome variable that captures fear of job replacement or hope of productivity gains. Recall that this estimation will be based on the group of respondents in Experiment T.

Expectations about policy preferences As we mentioned before, we lack firm expectations about the direction of β_1 in regression model (2). This is because this effect depends on the relative share of opinions regarding whether GenAI will produce job loss or productivity gains. This estimation will also be based on the group of respondents in Experiment T.

Expectations about perceptual mechanisms (H2a and H2b) To test the validity of the perceptual "job loss" and "productivity gains" mechanism, we consider regression model (3) which, recall, is based on respondents in Experiment TI. The expectations laid out in page 9 imply the following combinations of parameter estimates: $\beta_4 < 0$ (H2a) and $\beta_5 > 0$ (H2b).¹⁸

¹⁸Recall that the indicators for the outcome variables are coded to be positive, implying more support for defensive policies, less support for active policies, and more support for passive policies.

6.1 Additional Analyses

Follow-up survey analysis We supplement our analysis by conducting a robustness check using data from our one-week follow-up survey. We first run the following regressions

$$Y_i^{\text{Follow-up}} = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{AI Task}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_i \quad (4)$$

in which $Y_i^{\text{Follow-up}}$ represents participant i 's responses collected from the follow-up survey, including attitudes towards AI-related policies as well as expectations on AI's impact on the economy.

Additionally, as we track participant's knowledge about AI-related news, we also investigate the one-week impact of our intervention on participants' understanding about AI or AI-related knowledge. To do this, we estimate the following regression:

$$K_i^{\text{Follow-up}} = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{AI Task}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_i \quad (5)$$

in which $K_i^{\text{Follow-up}}$ represents participant i 's knowledge about AI-related news collected in our follow-up surveys.

AI tools and participants' performance In our experiment, each participant is asked to complete a complex multi-step task, either with or without the assistance of AI tools. By recording the participants' responses to each step, we aim to explore how the introduction of AI tools impacts their ability to manage complex job-related tasks. To deepen our analysis, we estimate the following regression model:

$$Q_{ij} = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{AI Task}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_{ij} \quad (6)$$

in which Q_{ij} is a set of variables that capture the performance of participant i in task j . Performance indicators include length of the text submitted for the task, time taken to complete each task, and whether answers are correct.

6.2 Robustness checks

We perform a battery of robustness checks based on the empirical specifications stated in Section 6.

Adjustment for engagement rate One of the key concerns that may affect identification and estimated effects is the variation in participant engagement levels. Although our

experiment is carefully designed with incentives to encourage thorough participation, we nonetheless conduct a series of robustness checks to account for potential issues related to participant engagement. This is because some individuals may complete the tasks with diligence and attention to detail, while others may exhibit lower levels of care, leading to inconsistent exposure to the different tools (that is, AI versus non-AI tools). To mitigate this issue, we collect detailed engagement metrics, including time spent on each task, length and complexity of responses (measured by word and sentence counts), and frequency and duration of interactions within each intervention arm of our main experiment. These variables serve as proxies for participant engagement. We then perform a sub-sample analysis focusing on participants who completed the tasks correctly, allowing us to assess the robustness of our findings under conditions of higher engagement.

Country-specific subsamples We plan to conduct our experiment in five Latin American countries. The populations of these countries have different cultures and economic conditions, and participants could therefore respond differently and in unforeseen ways to the experimental interventions. Therefore, our results might be driven by samples from a specific country, or vary dramatically across different countries. To eliminate these possibilities, we will re-estimate all regression models based on each country’s subsample. The purpose is to identify “outlier country samples”, not to estimate within-country effects, for which national samples are under-powered.

Attrition adjustments Although we do not anticipate any differences in the fundamental characteristics of participants between our baseline and final survey samples, we will adjust our estimates to take into account the possibility of differential attrition across treatment arms using the methodologies of Lee (2009).

6.3 Heterogeneous effects

We consider potential heterogeneity by focusing on the following moderators: *(i)* gender; *(ii)* education level; *(iii)* personal income; *(iv)* left-right self-placement; *(v)* religious affiliation; and *(vi)* their pre-treatment knowledge and experience of GenAI. We see this as a more limited descriptive step to analyze whether individuals in these social groups differentially perceive threats and opportunities from AI.

We devote some additional space to a test of heterogeneous effects by employment sector. As individuals may form perceptions about AI differently based on their personal

experiences, we also estimate the following regressions:

$$P_i = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{AI Task}_i} + \tilde{\beta}_2 1_{\text{AI Task}_i} \times 1_{\text{High AI Intensity}_i} + \tilde{\beta}_3 1_{\text{High AI Intensity}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_i \quad (7)$$

$$Y_i = \tilde{\beta}_0 + \tilde{\beta}_1 1_{\text{AI Task}_i} + \tilde{\beta}_2 1_{\text{AI Task}_i} \times 1_{\text{High AI Intensity}_i} + \tilde{\beta}_3 1_{\text{High AI Intensity}_i} + X_i \tilde{\gamma} + \tilde{\varepsilon}_i \quad (8)$$

Here, $1_{\text{High AI Intensity}_i}$ indicates whether individual i has a job with multiple tasks that can be easily done by generative AI. Since one of our hypotheses is that individuals' experiences in the labor market may shape their perceptions of AI's impact, we primarily focus on the coefficient of the interaction term between interaction with AI to accomplish tasks and working in occupations that are potentially substitutable by AI, that is, the estimated value of $\tilde{\beta}_2$ in Equation 8. Our measure of high AI intensity combines participants' self-reported occupations with an index developed by Felten, Raj and Seamans (2021). In the survey, participants select their current job from a list of sectors and professions aligned with the O*NET¹⁹ classification system. We then match each respondent's occupation to the dataset of Felten, Raj and Seamans (2021).

6.4 Limitations and challenges

The risks associated with this study are minimal. The primary concern with implementing the field experiment is the potential for delays due to communication issues with survey companies. However, this is not a significant issue, as our experiment does not depend on any specific future event. Conditional on approval of our pre-registered design, we think it is very likely that our experiment will be completed before December 2025.

Our experiment poses minimal risks to participants. We have two types of interventions: informing participants about the general risks and opportunities brought about by AI, and incentivizing them to use AI tools to solve a series of tasks. We *do not* disseminate fake news or misinformation that could distort participants' understanding of the current economic situation in their countries. We also offer training on how to use AI tools to some participants, which should be beneficial, or at least harmless, in their daily lives. Our interventions adhere to the traditional rules of field experiments and have been approved by our institutional review boards to address potential ethical concerns.

We address potential challenges such as attrition and non-compliance in Section 5.3. We have carefully designed our experiment to ensure that our interventions and tasks are easy for participants to understand, thus minimizing the likelihood that such issues might arise during our experiment.

¹⁹<http://onetonline.org/>

7 Policy implications

This study provides timely evidence about individuals' beliefs and attitudes about the introduction of AI, a technology that continues to evolve at neck-breaking speed, thereby assisting policymakers in anticipating and addressing its future societal impacts. Understanding people's perceptions of AI and its potential consequences is crucial for designing policies that maximize productivity benefits while minimizing adverse effects. In particular, this study sheds light on how individuals' attitudes toward AI and related policy preferences are influenced by direct interaction with AI tools to solve complex tasks in work environments and the perceptions that they form regarding the potential negative and positive effects of this technology. These are channels that are distinct from, and may not fully reflect, AI's actual labor market outcomes. These implications are particularly relevant for middle-income developing countries, like the ones we analyze, that remain passive recipients of new technologies and possess limited institutional capacity to address AI-related challenges. More broadly, our findings will inform strategies to reduce societal disruptions historically associated with the widespread adoption of transformational technologies, such as social unrest (Caprettini and Voth, 2020).

Among the multiple interventions that policymakers in these countries could envision, we anticipate a heavier emphasis on policies that promote equitable access to AI, on the one hand, and policies that promote protection from potential employment loss, on the other. To promote equitable access to AI, governments could fund AI-training programs in public high schools and colleges, guaranteeing that lower- and middle-income students can develop AI-related skills at the same pace as students in higher-income private institutions. Policymakers could also subsidize internet access and computer infrastructure for small businesses, which are unlikely to successfully transform on their own to take advantage of the AI revolution. Policymakers could also promote open-source policies to prevent the creation of digital monopolies that would thwart access to AI tools.

Alongside policies that promote the use of AI, governments could partner with private actors to create re-skilling initiatives, offer financial support to workers displaced by AI and automation, and tax companies that substitute jobs for AI so as to slow the pace of the transition and in order to fund unemployed individuals. Farther away from direct interventions that affect specific firms and workers, policymakers could consider broad policies that reduce the number of working hours among current employees, and policies like a universal basic income that guarantee access to steady income streams regardless of job status. These policies could diminish the potentially negative effects of a full

AI transition, but they can only be successful if we can anticipate the likely reaction of individuals, which is one of the purposes of our study.

We emphasize that our potential conclusions apply primarily to individuals in developing countries who have internet access and limited prior exposure to AI.²⁰ We remain cautious about extending our findings to major AI-producing countries, such as the United States and China, where baseline knowledge of AI may be higher and governments possess greater institutional capacity and more extensive policy instruments to proactively address AI-related challenges.

²⁰Although we rely on online-panel samples, we believe the results are broadly applicable to the general populations of our targeted regions, as online-panel samples have been shown to produce response patterns and characteristics similar to those in offline probability-based samples (Coppock and McClellan, 2019; Grewenig et al., 2023; Haaland, Roth and Wohlfart, 2023).

References

- Acemoglu, Daron. 2002a. "Directed technical change." *The Review of Economic Studies* 69(4):781–809.
- Acemoglu, Daron. 2002b. "Technical change, inequality, and the labor market." *Journal of Economic Literature* 40(1):7–72.
- Acemoglu, Daron. 2021. "Harms of AI". Technical report National Bureau of Economic Research.
- Acemoglu, Daron. 2024. "The Simple Macroeconomics of AI". Technical report National Bureau of Economic Research.
- Acemoglu, Daron, Asuman Ozdaglar and James Siderius. 2024. "A model of online misinformation." *Review of Economic Studies* 91(6):3117–3150.
- Acemoglu, Daron and James A Robinson. 2012. "Why Nations Fail" *New York Review of Books* 59(13):85–86.
- Acemoglu, Daron and Pascual Restrepo. 2018. "The race between man and machine: Implications of technology for growth, factor shares, and employment." *American Economic Review* 108(6):1488–1542.
- Acemoglu, Daron and Pascual Restrepo. 2019. "Automation and new tasks: How technology displaces and reinstates labor." *Journal of Economic Perspectives* 33(2):3–30.
- Acemoglu, Daron and Pascual Restrepo. 2020. "Robots and jobs: Evidence from US labor markets." *Journal of Political Economy* 128(6):2188–2244.
- Anelli, Massimo, Italo Colantone and Piero Stanig. 2019. We were the robots: Automation and voting behavior in western europe. Technical report CReAM Discussion Paper Series.
- Anelli, Massimo, Italo Colantone and Piero Stanig. 2021. "Individual vulnerability to industrial robot adoption increases support for the radical right." *Proceedings of the National Academy of Sciences* 118(47):e2111611118.
- Autor, David, David Dorn, Gordon Hanson and Kaveh Majlesi. 2020. "Importing political polarization? The electoral consequences of rising trade exposure." *American Economic Review* 110(10):3139–3183.

Autor, David H, Frank Levy and Richard J Murnane. 2003. "The skill content of recent technological change: An empirical exploration." *The Quarterly Journal of Economics* 118(4):1279–1333.

Benítez-Rueda, M. and E. Parrado. 2024. "Mirror, Mirror on the Wall: Which Jobs Will AI Replace After All? A New Index of Occupational Exposure." *Econstor* .
URL: <https://www.econstor.eu/bitstream/10419/302209/1/1900545608.pdf>

Bordalo, Pedro, Giovanni Burro, Katherine Coffman, Nicola Gennaioli and Andrei Shleifer. 2025. "Imagining the future: memory, simulation, and beliefs." *Review of Economic Studies* 92(3):1532–1563.

Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon and Andrei Shleifer. 2025. "How people use statistics." *Review of Economic Studies* p. rdaf022.

Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon and Andrei Shleifer. 2023. "Memory and probability." *The Quarterly Journal of Economics* 138(1):265–311.

Brynjolfsson, Erik, Danielle Li and Lindsey Raymond. 2025. "Generative AI at work." *The Quarterly Journal of Economics* 140(2):889–942.

Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg et al. 2023. "Sparks of artificial general intelligence: Early experiments with gpt-4." *arXiv preprint arXiv:2303.12712* .

Campello, Daniela and Francisco Urdinez. 2021. "Voter and legislator responses to localized trade shocks from China in Brazil." *Comparative Political Studies* 54(7):1131–1162.

Caprettini, Bruno and Hans-Joachim Voth. 2020. "Rage against the machines: Labor-saving technology and unrest in industrializing England." *American Economic Review: Insights* 2(3):305–320.

Card, David. 2009. "Immigration and inequality." *American Economic Review* 99(2):1–21.

Coffman, Katherine, Manuela R Collis and Leena Kulkarni. 2024. "Stereotypes and belief updating." *Journal of the European Economic Association* 22(3):1011–1054.

Coppock, Alexander and Oliver A McClellan. 2019. "Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents." *Research & politics* 6(1):2053168018822174.

Di Battista, Attilio, Sam Grayling, Elselot Hasselaar, Till Leopold, Ricky Li, Mark Rayner and Saadia Zahidi. 2023. Future of jobs report 2023. In *World Economic Forum, Geneva, Switzerland*. <https://www.weforum.org/reports/the-future-of-jobs-report-2023>.

Eloundou, Tyna, Sam Manning, Pamela Mishkin and Daniel Rock. 2023. "Gpts are gpts: An early look at the labor market impact potential of large language models." *arXiv preprint arXiv:2303.10130* .

Felten, Ed, Manav Raj and Robert Seamans. 2023. "How will language modelers like chatgpt affect occupations and industries?" *arXiv preprint arXiv:2303.01157* .

Felten, Edward, Manav Raj and Robert Channing Seamans. 2019. The effect of Artificial Intelligence on human labor: An ability-based approach. In *Academy of Management Proceedings*. Vol. 2019 Academy of Management Briarcliff Manor, NY 10510 p. 15784.

Felten, Edward, Manav Raj and Robert Seamans. 2021. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." *Strategic Management Journal* 42(12):2195–2217.

Fowler, Anthony, Gregory A. Huber, Rongbo Jin and Lilla V. Orr. 2024a. "Why Are You a Democrat? Studying the Origins of Party Identification and Partisan Animosity with Open-ended Survey Questions." *unpublished manuscript* .

Fowler, Anthony, Gregory A. Huber, Rongbo Jin and Lilla V. Orr. 2024b. "Why Are You a Democrat? Studying the Origins of Party Identification and Partisan Animosity with Open-ended Survey Questions." *unpublished manuscript* .

Fowler, Anthony et al. 2020. "Partisan intoxication or policy voting?" *Quarterly Journal of Political Science* 15(2):141–179.

Gallego, Aina, Alexander Kuo, Dulce Manzano and José Fernández-Albertos. 2022. "Technological risk and policy preferences." *Comparative Political Studies* 55(1):60–92.

Gallego, Aina and Thomas Kurer. 2022. "Automation, digitalization, and artificial intelligence in the workplace: implications for political behavior." *Annual Review of Political Science* 25:463–484.

Gmyrek, P., H. Winkler and S. Garganta. 2024. "Buffer or Bottleneck? Employment Exposure to Generative AI and the Digital Divide in Latin America." *J. Auditag-ec* . URL: <https://jauditag-ec.com/BM-IA.pdf>

- Gong, Xiaodong, Arthur Van Soest and Elizabeth Villagomez. 2004. "Mobility in the urban labor market: a panel data analysis for Mexico." *Economic Development and Cultural Change* 53(1):1–36.
- Graeber, Thomas, Christopher Roth and Florian Zimmermann. 2024. "Stories, statistics, and memory." *The Quarterly Journal of Economics* 139(4):2181–2225.
- Grewenig, Elisabeth, Philipp Lergetporer, Lisa Simon, Katharina Werner and Ludger Woessmann. 2023. "Can internet surveys represent the entire population? A practitioners' analysis." *European Journal of Political Economy* 78:102382.
- Guriev, Sergei and Elias Papaioannou. 2022. "The political economy of populism." *Journal of Economic Literature* 60(3):753–832.
- Gyöngyösi, Győző and Emil Verner. 2022. "Financial crisis, creditor-debtor conflict, and populism." *The Journal of Finance* 77(4):2471–2523.
- Haaland, Ingar, Christopher Roth and Johannes Wohlfart. 2023. "Designing information provision experiments." *Journal of Economic Literature* 61(1):3–40.
- Heckman, James J and Carmen Pagés. 2000. "The cost of job security regulation: Evidence from Latin American labor markets.".
- Hémous, David and Morten Olsen. 2021. "Directed technical change in labor and environmental economics." *Annual Review of Economics* 13:571–597.
- Holland, Alisha C. 2018. "Diminished expectations: Redistributive preferences in truncated welfare states." *World Politics* 70(4):555–594.
- Hui, Xiang, Oren Reshef and Luofeng Zhou. 2024. "The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market." *Organization Science* 35(6):1977–1989.
- IPSOS. 2024. *Public opinion survey on AI understanding in Latin America*. Ipsos Group S.A.
- Jiang, Zhengyang, Hongqi Liu, Cameron Peng and Hongjun Yan. 2024. Investor memory and biased beliefs: Evidence from the field. Technical report National Bureau of Economic Research.
- Karunakaran, Arvind, Sarah Lebovitz, Devesh Narayanan and Hatim A Rahman. 2025. "Artificial Intelligence at Work: An Integrative Perspective on the Impact of AI on Workplace Inequality." *Academy of Management Annals* (ja):annals–2023.

Latinobarómetro Corporation. 2025. "Annual Reports." <https://www.latinobarometro.org/latContents.jsp?CMSID=InformesAnuales&Idioma=0>. Accessed: 2025-03-29.

Lee, David S. 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects." *Review of Economic Studies* 76(3):1071–1102.

Malmendier, Ulrike and Stefan Nagel. 2011. "Depression babies: Do macroeconomic experiences affect risk taking?" *The Quarterly Journal of Economics* 126(1):373–416.

Milner, Helen V. 2021. "Voting for populism in Europe: Globalization, technological change, and the extreme right." *Comparative Political Studies* 54(13):2286–2320.

Molinder, Jakob, Tobias Karlsson and Kerstin Enflo. 2021. "More power to the people: electricity adoption, technological change, and labor conflict." *The Journal of Economic History* 81(2):481–512.

Mummolo, Jonathan and Erik Peterson. 2019. "Demand effects in survey experiments: An empirical assessment." *American Political Science Review* 113(2):517–529.

Nettel, Pablo Fuentes, Emma Hankins, Richard Stirling, Giulia Cirri, Gonzalo Grau, Sulamaan Rahim and Eddie Crampton. 2024. Government AI Readiness Index 2024. Report Oxford Insights.

URL: <https://oxfordinsights.com/ai-readiness/ai-readiness-index/>

Noy, Shakked and Whitney Zhang. 2023. "Experimental evidence on the productivity effects of generative artificial intelligence." *Science* 381(6654):187–192.

Page, Benjamin I and Richard A Brody. 1972. "Policy voting and the electoral process: The Vietnam War issue." *American Political Science Review* 66(3):979–995.

Rodrik, Dani. 2021. "Why does globalization fuel populism? Economics, culture, and the rise of right-wing populism." *Annual Review of Economics* 13(1):133–170.

Stantcheva, Stefanie. 2023. "How to run surveys: A guide to creating your own identifying variation and revealing the invisible." *Annual Review of Economics* 15(1):205–234.

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

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts and Siddharth Suri. 2025. "Working with AI: Measuring the Occupational Implications of Generative AI." *arXiv preprint arXiv:2507.07935* .

Wu, Nicole. 2022. "Misattributed blame? Attitudes toward globalization in the age of automation." *Political Science Research and Methods* 10(3):470–487.

Appendix A Intervention details

A.1 Statements in AI information treatment

Information about Productivity Gain from AI Adoption Artificial intelligence (AI) is rapidly transforming how individuals work by automating routine tasks and allowing people to focus on what they do best. According to a Goldman Sachs report, AI's disruptive potential is set to reshape the labor market, and generative AI is already approaching levels of performance nearly indistinguishable from human output. By handling repetitive and time-consuming activities—from data entry and scheduling to drafting emails—AI frees up valuable cognitive resources, enabling individuals to engage in more creative, strategic, and high-value work. For example, research from MIT Sloan demonstrates that generative AI can boost worker productivity by nearly 40% in controlled experiments, while studies by the Nielsen Norman Group have shown that AI assistance can increase throughput by up to 66% in practical tasks.

In the future, by assisting people with routine and repetitive tasks, AI could help them focus on activities where they have comparative advantages, thus improving their productivity. According to research by economist Daron Acemoglu, the 2024 Nobel Prize in Economics laureate, with the development of AI, a large number of new jobs will be created through this technological progress.

Recently, a team of economists and computer scientists from Princeton University and New York University calculated an AI-relatedness score for each job in the U.S. labor market. A higher AI-relatedness score indicates that AI can be more helpful in that job. Below is a table with three examples representing high, medium, and low AI-relatedness scores.

Job	AI Score	Rating	Examples of tasks in which AI could be helpful
Telemarketers	1.93	(High)	Cold Calling, Scripted Sales Interactions, Follow-Up Scheduling
Pharmacy Technicians	-0.02	(Medium)	Prescription Processing, Medication Dispensing, Drug Interaction Checking
Laborers and Freight, Stock, and Material Movers	-1.50	(Low)	Most tasks cannot be performed by current AI

Information about Job Losses Caused by AI Artificial intelligence (AI) is rapidly transforming how individuals work by automating routine tasks—but this transformation is also accelerating the displacement of human workers. According to a Goldman Sachs report, AI's disruptive potential could render up to 25% of current work obsolete, potentially affecting as many as 300 million full-time jobs globally by 2030. By taking over repetitive and time-consuming activities—from data entry and scheduling to drafting emails—AI not only cuts costs but also removes the need for human intervention, threatening vast numbers of jobs and eroding workers' livelihood opportunities. For example, according to the OECD, adoption of AI technologies could result in approximately 14% of the labor force being displaced, with particularly high vulnerability in labor-intensive sectors such as services, retail, and logistics. Similarly, McKinsey's recent analysis estimates that nearly 12 million US workers in shrinking occupations may be compelled to transition out of their current roles by 2030 due to generative AI, underscoring a significant negative impact on job security.

In the future, by replacing humans in tasks that AI can perform at lower costs, AI could lead to significant unemployment and wage reductions for many people. According to research by

economist Daron Acemoglu, the 2024 Nobel Prize in Economics laureate, the development of AI will not reduce economic inequality and is even likely to have a negative impact on the real earnings of participants in the labor markets.

Recently, a team of economists and computer scientists from Princeton University and New York University calculated an AI-relatedness score for each job in the U.S. labor market. A higher AI-relatedness score indicates a greater likelihood that AI will replace the tasks involved in that job. Below is a table with three examples representing high, medium, and low AI-relatedness scores.

Job	AI Score	Rating	Examples of tasks that AI could replace
Telemarketers	1.93	(High)	Cold Calling, Scripted Sales Interactions, Follow-Up Scheduling
Pharmacy Technicians	-0.02	(Medium)	Prescription Processing, Medication Dispensing, Drug Interaction Checking
Laborers and Freight, Stock, and Material Movers	-1.50	(Low)	Most of the tasks cannot be performed by current AI

Irrelevant information about the entertainment industry The entertainment industry has become a part of everyday life. In recent years, the range of entertainment options has grown, including traditional forms like theater and music concerts, alongside newer options such as video games and streaming platforms. These different formats cater to varied preferences among global audiences. With the advent of digital technology, the industry has rapidly evolved, offering more interactive and personalized experiences that transcend traditional boundaries. Additionally, the integration of social media has enabled audiences to engage more directly with their favorite content, fostering a dynamic community that continuously shapes the way entertainment is consumed.

A report from cultural studies researchers notes that the entertainment industry continues to generate substantive profits. In 2024, the film and music sectors drew attention, with Hollywood's latest superhero movie, *Guardians of Tomorrow*, seeing significant box office performances. Meanwhile, international music acts like the South Korean group BTS maintained a presence on global charts, reflecting ongoing shifts in musical trends.

Cultural festivals and events are also held in various regions, such as Brazil's Carnival and Japan's Anime Expo. These events serve as occasions for people to experience aspects of different cultures. According to studies, such activities are often seen as opportunities for cultural engagement and shared experiences, which promote broader empathy among people with different cultural backgrounds.

A.2 Benchmark task

The first part of the questionnaire asks you to consider a hypothetical scenario and to complete two multi-step office tasks. These timed tasks require concentration and will take approximately 20 to 25 minutes. Some questions are open-ended and do not have a single correct answer. At the end of the study, we will raffle five Amazon gift cards, each valued at \$20, among respondents who complete both the tasks and the follow-up survey with high quality. **In the first task, we will provide you with a calculator in the bottom-left corner of your screen.**

This is the hypothetical scenario: You are a senior operations manager at GlobalTech Solutions, a multinational technology firm that assembles computers in Mexico. Your company is preparing to launch a new computer brand, scheduled for September 30, 2025. You are working with logistics partners who ship critical electronic components—especially computer processors—to your assembly plant.

Recently, you received the following email:

An email message displayed here (following is the email)

Dear Team,

We are writing to inform you about some urgent shipment issues:

Phoenix batch might be delayed 1–2 days due to inspection.

Austin shipment showed only 150 units instead of 1,500.

San José shipment is expected to arrive in 5 days.

We urgently need 1,500 working computer processors to keep production on track.

Best regards,

QuickShip Logistics

In addition, we have provided the Supplier Information Table below, which we believe may be helpful for your consideration.

Location	Capacity (available units)	Failure rate (%)	Transport cost (\$/unit)	Shipment length (days)	Price (\$/unit)
Austin, TX	700	1.0	10	2	270
Phoenix, AZ	800	2.0	12	3	260

You will have a limited amount of time to answer these questions. In each page, this hypothetical scenario will be reproduced for your convenience.

[Task 1 question] Which of the three issues do you consider to be most critical for production?

a = Short delay in Phoenix

b = Austin missing 1,350 units

c = Longer shipping time in San Jose

[Task 2 question] Suppose you order all available units from Austin and Phoenix. How many functional processors can you expect after accounting for the failure rate? Formula: Expected = Units \times ((100 – Failure Rate) / 100). Example: 800 units at 1.0% Failure Rate \rightarrow 800 \times ((100 – 1)/100) = 792.

[Task 3 question] Assume you are David, the Operations Manager at GlobalTech Solutions. Draft a short professional email (no more than 100 words) to QuickShip Logistics requesting an explanation for the significant delivery shortage (only 150 processor units delivered instead of the expected 15,000).

Recipient: Ms. Rachel Thornton
Email: rthornton@quickship.com

Remember: Your email should use professional business English, include a subject line, and be under 100 words

Draft your email below:

A.3 Task questions in AI exposure treatment

Task description The next part of the questionnaire asks you to consider a hypothetical scenario and requires you to perform four connected tasks. These tasks require concentration, and will take you anywhere from 20 to 25 minutes. The tasks are open-ended, there are no right or wrong answers. At the end of our study, we will raffle two prizes worth XX Korus among respondents that have finished the tasks. (For treatment group: We invite you to use ChatGPT to address these tasks. Below, we will ask you to see a brief video that shows you how to interact with ChatGPT and, especially, how to formulate effective prompts for ChatGPT.)

A video is shown here to help participants understand how to effectively prompt AI

Now consider the following hypothetical scenario: You are a senior operations manager at GlobalTech Solutions, a multinational technology firm that assembles computers in Mexico. Your company is preparing for the launch of a new computer brand, and you are overseeing the logistics and supply chain operations. You are working with QuickShip Logistics, an external logistics partner responsible for shipping critical electronic components — including computer chips — from multiple suppliers around the world to your assembly plants in Mexico. These shipments are essential for final assembly before the product launch on April 15, 2025.

[Task 1 prompt] Your boss has just forwarded to you the following email, which is poorly written and unclear. He is concerned that GlobalTech Solutions will not be able to launch the product in time, but he does not understand which of the following failures from QuickShip Logistics is the one that he needs to address. Your first task is to identify this crucial problem:

An email message displayed here (following is the email)

Hey,

Hope ur doing well. Need to bring up some big problems with shipments & delays that cud rlly mess with our operations if not fixed fast. A few of these are minor concerns, but one of them is a big issue since you want to build new computers!

Shanghai Shipment – Supposed to arrive Feb 25, but stuck bcz of new import docs. No clear date from customs, making it hard to plan production of new cars.

LA Shipment – Early March delivery, but now held for “random inspection” by US customs. No timeline = big logistics mess. We could end up losing the big contract with that X-ray company.

Houston Shipment – Jan 28 audit found only 150 computer processors instead of 15000. This is messing up assembly & causing admin headaches. Oops!!

Shenzhen Supplier Delay – Shutdown for maintenance. Was due March 10, but now won’t confirm til March 9. Not much time to find freezer bags elsewhere.

Other Issues – Bad communication on shipments, handling probs, not enough updates from ur team, supply chain bottlenecks in all major ports.

We need to talk ASAP to fix this & figure out backup plans. Pls let me know ur thoughts soon.

[Task 1 question] Which of the issues mentioned in the email is most likely to disrupt the launch of the new computer brand? [Open-ended answer]

After task has been completed, allow the respondent to advance to the next screen

[Task 2 question] Your boss now understands that the problem is the lack of computer processors. He provides you with a list of alternative suppliers. For each supplier, you have access to information on total factory capacity, the failure rate of each processor, the expected cost of transportation (per unit), the expected delay in receiving a shipment, and the actual price of the processor (per unit).

Location	Capacity (available units)	Failure rate (%)	Transport cost (\$/unit)	Shipment length (days)	Price (\$/unit)
San José, Costa Rica	5,000	1.2	15	5	250
Shenzhen, China	9,000	2.5	30	12	220
Suzhou, China	8,500	2.0	28	11	230
Austin, Texas	6,000	0.8	10	2	270
Phoenix, Arizona	7,500	1.0	12	3	260

If you ordered the full capacity from all suppliers, calculate the total purchase cost of all processors, ignoring any transportation costs, and round your result to the nearest whole number. Please consider the failure rate of this process.

After task has been completed, allow the respondent to advance to the next screen

[Task 3 question] Your boss now asks you to contact QuickShip Logistics directly. He wants you to draft a professional email that: (i) Explains that you are very concerned about the processor

shortage (only 150 units delivered instead of 15,000). (ii) Requests an immediate clarification about how this mistake happened. (iii) Asks QuickShip Logistics to propose a concrete plan to fix the problem (e.g., reshipping, alternative suppliers, or compensation). (iv) Emphasizes the urgency of the situation, since the new product launch is scheduled for September 2025.

Recipient: QuickShip Logistics Customer Support

Email: support@quickship.com

Remember: Your email should use professional business English, include a subject line, and be under 100 words


Draft your email below:

A.4 AI interface

Figure A1: AI Interface

Question 1

Your boss has just forwarded the following email to you. It is poorly written and lacks clarity. He is concerned that GlobalTech Solutions may not be able to launch the product on time, but he does not know which of the following failures from QuickShip Logistics is the one that requires immediate attention. Your first task is to



AI Assistant

Hello! I'm your AI assistant, powered by OpenAI's GPT-5. I've gone through today's materials—please don't worry about copying the questions again. How can I be of assistance?

12:03 AM

Enter here Submit

Dear Sirs, we are doing well. Need to bring up some big problems with shipments & delays that could really hurt our operations if not fixed fast. A few of these are minor concerns, but one of them is a major issue since you want to build new computers!

Shanghai Shipment – Supposed to arrive Jul 25, but stuck bcz of new import docs. No clear info from customs, making it hard to plan production of new cars.

Shenzhen Shipment – Early August delivery, but now held for "random inspection" by US customs. No info = big logistics mess. We could end up losing the big contract with that X-ray company.

Shenzhen Supplier Delay – Jun 28 audit found only 150 computer processors instead of 15000. This is causing assembly and causing admin headaches. Oops!!

Shenzhen Supplier Issues – Bad communication on shipments, handling probs, not enough updates from our supplier, supply chain bottlenecks in all major ports.

We need to talk ASAP to fix this and figure out backup plans. Pls let me know ur thoughts soon.

Which issue mentioned in the email is most likely to disrupt the launch of the new computer brand?

✗ close assistant

For the convenience of our respondents in the treatment group, we incorporate the AI ChatBot with access to the GPT-5 model, as Figure A1 shows,²¹ in our survey platform.

²¹The actual interface will be in Portuguese or Spanish.

A.5 Calculator and task assistant interface (control group)

Figure A2: Calculator and Task Assistant (Control Group)

The screenshot shows a user interface for a task assistant. At the top, a header reads "Question 1". Below it is a text block: "Your boss has just forwarded the following email to you. It is poorly written and lacks clarity. He is concerned that GlobalTech Solutions may not be able to launch the product on time, but he does not know which of the following failures from QuickShip Logistics is the one that requires immediate attention. Your first task is to crucial problem:"

Task Hints

Welcome! This is a Task Hints. Enter a keyword to get the corresponding task hint: question1, question2, question3. Type help to see all available keywords and rules; you can enter multiple keywords at once (e.g., question1, question3).

11:51 PM

The main content area contains several bullet points of text:

- Shanghai Shipment – Supposed to arrive Jul 25, but stuck bcz of new import docs. No clear from customs, making it hard to plan production of new cars.
- Shipment – Early August delivery, but now held for "random inspection" by US customs. No line = big logistics mess. We could end up losing the big contract with that X-ray company.
- Shipment – Jun 28 audit found only 150 computer processors instead of 15000. This is sing up assembly and causing admin headaches. Oops!
- Shenzhen Supplier Delay – Shutdown for maintenance. Was due August 10, but now won't firm til August 9. Not much time to find freezer bags elsewhere.
- er Issues – Bad communication on shipments, handling probs, not enough updates from ur, supply chain bottlenecks in all major ports.

We need to talk ASAP to fix this and figure out backup plans. Pls let me know ur thoughts soon.

To mitigate potential engagement effects in our experiment, we offer participants access to a calculator and a chatbot called "Task Hints": tools designed to support task completion by providing basic utilities and guidance. While these features offer helpful hints, the rule-based chatbot cannot provide direct solutions to the task questions and will respond only when participants type the correct keywords. Participants using these tools are still required to perform fundamental calculations, conduct statistical analyses, and engage in reading and summarization to arrive at the correct task solutions.

A.6 Additional information provision experiment

Experiment I: Information-Only Treatment We perform an ancillary experiment on 3,000 individuals designed to assess whether exposure to information about AI, without any task interactions, can alter perceptions and, in turn, shift policy preferences. Although analyzing the

effects of different types of information is not the primary focus of our study — rather, we see the information-framing intervention as an opportunity to test the perceptual mechanisms that we believe mediate the relation between GenAI interaction and policy preferences — this sub-sample serves as a useful benchmark for Experiment TI. Participants are randomly assigned to one of three conditions: the control sample (30%, approximately 900 individuals), which receives information unrelated to AI; a sample (35%, approximately 1,050 individuals) that receives information about GenAI's potential job displacement effects; and another sample (35%, approximately 1,050 individuals) that receives information about GenAI's potential productivity gains.

Online Appendix

Please click on the following link to see the full survey:

https://drive.google.com/file/d/1XVJJJawXLtR05QIRYtrp_mfH1cuZeqmY/view?usp=sharing