

Collective action from labor in the age of AI: Theory and evidence from the advent of ChatGPT*

Carlos Felipe Balcázar[†]

July 4, 2026

Abstract

Can AI increase collective action from workers? To study this question I exploit a major event-shock: the advent of ChatGPT. I posit that the AI innovation incentivizes collective action via collective bargaining because workers anticipate a fall in their bargaining power after AI is implemented at the workplace. Hence collective action from labor should increase in response to the advent of ChatGPT by mobilizing preemptively. Using a differences-in-discontinuities design and data on collective bargaining agreements in the US, I find that workers are more likely to participate in collective bargaining agreements in response to the shock. Using the availability bias heuristic, I argue and show that the effect is driven by those previously exposed to advances in ICTs. These findings contribute to understanding of the political consequences of AI as force shaping collective action.

Words: 7324

*The comments of Michael Becher, Daniel Stegmüller, and participants at Yale's "New Challenges for International Cooperation: Automation, AI and Digital Trade" conference, NYU's International Relations seminar, ITAM's Political Economy seminar, "Navigating the New Frontier: The Political and Economic Implications of AI" at Caltech, are especially acknowledged.

[†]Department of Political Science and School of Public Policy, University College London, 29-31 Tavistock Square London, WC1H 9QU, UK. E-mail: c.balcazar@ucl.ac.uk.

1 Introduction

What are implications of new labor-replacing technologies for collective action? One of the most salient events in the past decades has been the advent of Generative Pretrained Transformers (GPTs) and its use in Artificial Intelligence (AI). These technologies have drawn the attention from policymakers and scholars alike because they represent technologies that could automate the tasks being performed by millions of workers. Consequently, some scholars have been concerned about the possibility that workers may oppose automation due to economic concerns and increase political instability—as a form of new Luddism (Frey, 2019). Nevertheless, there is no evidence of a “rage against the machines” (Zhang, 2022; Wu, 2022; Menon and Zhang, 2025).

Recently, however, we have observed increased worker mobilization in industries affected by these technologies, as the Hollywood strikes in 2023 exemplify: The Writers Guild of America and Screen Actors Guild-American Federation of Television and Radio Artists—representing more than 160 thou. workers in entertainment—went on strike over a labor dispute with the Alliance of Motion Picture and Television Producers for 148 days—on of the longest strikes in their history—to demand strict and harsh limits on using GPTs to replace workers, among other things. Workers in other industries exposed to AI such as health care, auto, technology, and other, are also organizing and striking not only in the US, but also around the world. These events echo anecdotal evidence of workers mobilizing in response to robot adoption starting in the 80s (Kennedy et al., 1982; Gil, 1986; Haapanala et al., 2023).

These events are taking place when unionization has drastically declined across the world, and evidence indicating the automation has weakened union power (Balcazar, 2023; Agnolin et al., 2025). The tech sector, in particular, is unionizing in response to AI despite exhibiting remarkably low levels of collective action in the past (Rothstein, 2022). This begs the question: does automation strengthen or weaken the political power of labor? I contribute to answer this question herein, but unlike previous work I focus on AI as opposed to robot adoption.

First of all, this question is inherently relevant for politics: Collective action from labor is of importance in the creation of both democratic and economic institutions. For instance, unions played a large role in pushing for democratic reforms and in creating welfare state institutions (Kaplan and Naidu, 2025). Unions act as grassroots movements that help voters, both union members and nonmembers, acquire costly political knowledge; they also coordinate voters and get them registered and then delivered to the polls, increasing the turnout of low- and middle-income individuals; they facilitate pooling resources from workers, which are often used for making campaign donations and for lobbying, to further the collective interests of workers; unions counterbalance private political interests, by advocating for redistributive policies in favor of workers. There is

essentially no social group that has the ability of coordinating workers' interests like unions do, both inside and outside the workplace (Dean, 2025).

Second, the impact of AI on collective action from labor remains unexplored in political science. Despite this, recent theoretical work on the impact of automation on collective action from workers indicate that automation can both increase and decrease this collective action (Balcazar, 2025; Becher and Stegmüller, 2025; Gingrich et al., 2025). Balcazar (2025) in particular theorizes that while innovations like AI can increase the level of distributive conflict between firms and workers because workers expect lower returns to work with the prospect of automation, capital deepening reduces workers' relative power vis-à-vis that of firms because their ability to reduce production by engaging in rent-seeking (e.g., strikes) instead of working, falls. Thus new labor-replacing technologies can incentivize collective action from workers, while their adoption reduces it.

I contribute to disentangling the question above by investigating the effect of AI on collective bargaining agreements in the US. In the U.S., the collective bargaining process involves substantial internal communication among union members to satisfy the legal requirement of a "sufficient community of interest," as encoded in the National Labor Relations Act of 1935. During this process, union members determine which issues to negotiate with employers and formulate the position of the bargaining unit; weak unions are less like to engage in collective bargaining (Becher and Stegmüller, 2021). Thus studying collective bargaining agreements is commensurate with the idea of studying collective action from labor. I also consider the role of strikes, which emerge when the bargaining process fails.

To tackle this endeavor I first propose a game theoretic model of collective action from workers using Rubinstein bargaining (Rubinstein, 1982). Unlike previous models, the Rubinstein bargaining adds strategic elements like the threat of disagreement, which hampers collective action, representing a more appropriate approach for the purpose herein. The model inquires about the impact of technological innovations (like AI) and their adoption on workers bargaining power. In it, I consider the market for automation and labor in reduced form, the endogenous decision to unionize, and the role of worker heterogeneity—which hampers unionization (Acemoglu et al., 2001).

I show that despite an innovation like AI may reduce the demand for labor, workers increase their level of collective action as technological innovation takes place and AI becomes more affordable for firms. They do so because they anticipate a fall in their bargaining power once AI is implemented—i.e., when the innovation is accompanied by capital deepening. This is more

likely to happen where the pre-existent level of capital deepening is high.¹ Further, the incentives for increased collective bargaining are attenuated if the labor force is heterogeneous or if there are frictions in the labor market, because the opportunity cost of collective action vis-à-vis doing productive activities increases. This mechanism shows that collective bargaining and strikes can increase in response to automation.

Second, I use the advent of ChatGPT in November 30th of 2022 in an innovative event study design to test the hypotheses above.² I focus the analysis on the US because of the wealth of high-frequency collective bargaining data that coincided with the event, which allows for a sound empirical exploration; the data comprises the universe of collective bargaining agreements in the US.

I use task routinization measure by Autor (2013) to measure the level of treatment uptake; the measure captures the implications of Information and Communication technologies (ICTs) on the demand for workers with different skills and characteristics prior to AI. I use this measure because workers should be more familiar with the potential consequences of AI if they have been previously exposed to changes in ICTs—this is known as *availability bias* (Tversky and Kahneman, 1973, 1974). This availability bias commensurates with the pre-existent level of capital deepening.

I measure the level of the AI shock event by computing a (properly) normalized level of interest in ChatGPT at the media-market level in the US, using data from Google trends. I show that this event has four advantages: i) We clearly observe a discontinuity in interest in ChatGPT after its release. This allows me to estimate the causal effect of the informational treatment. ii) We observe that lay-offs in the private sector coincide with the release of ChatGPT, making the event salient. iii) AI narratives emerged after the release of ChatGPT, when GPT4 was released in 2023. As a result, I find evidence of anticipation for the latter event and subsequent AI-technology releases, precluding the ability of using them. iv) Lay-offs fell after the shock because the market has been adjusting by generating new job opportunities (e.g., Aghion and Bunel 2024; Acemoglu 2025; Brynjolfsson et al. 2025), making this shock unique in its nature.

To address any concerns regarding the timing of the event I make use of differences-in-discontinuities design and optimal bandwidths around the event (Grembi et al., 2016; Eggers et al., 2018). I also perform manipulation and balance checks for reassurance.

I find that one standard deviation increase in exposure to ChatGPT is related to a statistically

¹This is consistent with (Becher and Stegmüller, 2025), but it relaxes the assumption of costly firm relocation they use.

²Although AI became commercially feasible and scalable after the transformer architecture debuted in 2017, it was not until the release of ChatGPT that this innovation generated an economic shock that pushed companies to quickly adopt AI to increase productivity and competitiveness.

significant effect of ChatGPT on collective bargaining agreements. In particular, the ChatGPT shock increases the number of collective bargaining agreements by 26.4 percentage points (or about one tenth of a standard deviation), and the likelihood of starting a new collective bargaining process by 3.1 percentage points. I also find that the effects of the advent of ChatGPT on collective action from labor are driven by workers in industries that were previously exposed to ICTs—consistent with the availability heuristic.

The findings herein contribute mainly to the following literatures: First, to the literature on the political economy of automation ([Anelli et al., 2018](#); [Thewissen and Rueda, 2019](#)), especially the one related to its effects on distributive conflict and collective action ([Balcazar, 2023](#); [Balcazar., 2025](#); [Becher and Stegmüller, 2025](#); [Agnolin et al., 2025](#)). Second, I also contribute to the literature about the effects of economic shocks on unionization ([Alhquist and Downey, 2021](#); [Becher and Stegmüller, 2023](#); [Owen and Park, 2025](#)). I demonstrate that the advent of AI manifests in increased distributive conflict between firms and workers via more collective bargaining agreements. This occurs because the AI innovation reduces workers’ opportunity cost from unionization and thus they “preemptively” start unionizing, as their bargaining power should fall once AI is implemented via capital deepening.

I also contribute to the literature by showing that the effects of AI on collective action are stronger in those whose industries where workers were more affected by ICTs in the past. The literature has found consistently that those who should be more threatened given current measures of exposure to AI, are not those that feel more threatened by it ([Menon and Zhang, 2024](#); [Anelli et al., 2025](#)), instead the labor market effects may depend on impacts by previous technological shocks (e.g., [Baccini et al. 2025](#)).

2 Theory

I first present the model without endogenous labor markets, frictions and heterogeneous workers to pin down the strategic incentives of workers and firms, and facilitate the exposition. Then I extend the model to include these aspects.

An employer (E) invests amount $r \geq 0$ in automation—like AI—and hires a mass $L > 0$ of workers to produce a generic good. Each worker, $i \in [0, L]$, receives a market wage $w > 0$ from the employer. The marginal cost of one unit of AI is $p > 0$. Changes to this quantity capture accessibility to the technology, thus $\Delta p < 0$ can occur as a result of an innovation that facilitates AI adoption. The production function $F(\cdot)$ satisfies $F(0) = 0$, it is increasing and concave in its inputs, and it is twice differentiable.

All markets are in equilibrium. The marginal returns to both workers and AI are their marginal productivities.

Workers pay a fee $c > 0$ to join their union, furthermore, workers display idiosyncratic tastes about belonging to a union. This captures the observation that the likelihood of belonging to a union is a function of workers' characteristics. The taste parameter is given by ε_i , where $\{\varepsilon_i\}_{i \in [0, L]}$ are distributed i.i.d. with mean zero and symmetric cumulative distribution function $G(\cdot)$, with full support on $\mathbb{R}$. Thus ε_i is a measure of political, ideological and other idiosyncratic differences between the worker and the union. Note that the existence of right-to-work laws and other frictions reducing the benefit of unionizing are also captured in reduced form by these parameters.

Each worker independently and simultaneously decides whether to join the labor union or not, and then the employer and the labor union bargain unionized workers' wages á-lá Rubinstein (Rubinstein, 1982). Unionized labor participates in production only when the employer and the union have reached an agreement, otherwise only non-unionized labor works. This captures the ability of workers to strike when collective bargaining fails.

Denote by w^* the wage of unionized workers and $n^* \in [0, L]$, their mass. Then: the employer's per-period payoff is

$$u_E = \begin{cases} F(r + L - n^*) - w \cdot (L - n^*) - r \cdot p & , \text{ no agreement} \\ F(r + L) - w \cdot (L - n^*) - w^* \cdot n^* - r \cdot p & , \text{ under agreement } w^* \end{cases}$$

We rule out $F(\cdot)$ linear to avoid corner solutions, and to remain consistent with the tasks based approach to studying automation (Acemoglu and Restrepo, 2020). Note also that marginal productivity of workers and machines is the same; that is, a unit of AI can replace a worker. This implies that AI occurs via capital deepening, because more AI reduces the labor share ceteris paribus. This implies in turn that in the absence of market frictions $p = w$ in equilibrium, otherwise $p \neq w$.

A worker's per-period payoff is

$$u_i = \begin{cases} w & , \text{ not unionized} \\ \varepsilon_i & , \text{ unionized and there is a strike} \\ w^* - c + \varepsilon_i & , \text{ unionized and the union reached an agreement } w^* \end{cases} ;$$

the labor union's pay-off is the average payoff of unionized workers:

$$u_U = \begin{cases} \int_0^L \varepsilon_i \cdot \mathbb{1}_i^U di & , \text{ a strike is ongoing} \\ w^* - c + \int_0^L \varepsilon_i \cdot \mathbb{1}_i^U di & , \text{ under agreement } w^* \end{cases} ,$$

where $\mathbb{1}_i^U$ takes the value of one if i is part of the union and zero otherwise.

The sequence of play is as follows:

- i) Each worker i decides to join the labor union or not.
- ii) The union bargains the wage of unionized workers á-la Rubinstein alternating offers.
 - When there is no agreement i does not work—if unionized—and receives its fee c .
 - When there is an agreement i works—if unionized—and gets w^* .
- iii) Production takes place, revenues are collected and wages are paid.

2.1 Bargaining wages and unionizing under the shadow of strikes

In the first period of bargaining, either the employer or the labor union start by proposing a wage w^* for the mass of unionized workers. After an agreement w^* is reached, it is implemented for the rest of the game. We assume that the employer discounts the future at a constant rate $\delta_E \in (0, 1)$, whereas workers and the labor union discount future payoffs at a rate $\delta_i \in (0, 1)$.

Proposition 1 characterizes the equilibrium in the sub-game when bargaining starts. We observe that union-workers' wages depend on the relative bargaining power of the labor union under the credible threat of a strike: $[F(r + L) - F(r + L - n^*)]/n$, conditional on the optimal level of AI adoption. The identity of the first mover during the bargaining process is inconsequential.

Proposition 1. *Assume a mass $n^* \in [0, L]$ of unionized workers, their wage is given by the following expression:*

$$w^* = \lambda \frac{F(r + L) - F(r + L - n^*)}{n^*} + (1 - \lambda)c, \quad \lambda \in \left\{ \frac{1 - \delta_e}{1 - \delta_i \delta_e}, \frac{1 - \delta_e}{1 - \delta_i \delta_e} \delta_i \right\} \subseteq [0, 1], \quad (1)$$

where the value of λ depends on who proposes first.

There are two elements that determines union-workers' wages: First, the opportunity cost from striking: $[F(r + L) - F(r + L - n^*)]/n$. This quantity corresponds to the relative bargaining power of the labor union under the credible threat of a strike. Second, the union fee, c . Since this fee is returned in full to the union worker during a strike, higher union fees reduce the benefits of an agreement vis-à-vis those of a strike.

A worker unionizes if and only if the cost of opportunity of doing is lower than the benefit: $w^* - c + \varepsilon_i > \min\{p, w\}$.³ By the law of large numbers, the fraction of unionized workers is $G(w^* - \min\{p, w\} - c)$ (Proposition 2).

Proposition 2. *Given the equilibrium wage w^* , the mass of unionized workers, $n^* \in [0, 1]$, satisfies the following equality:*

$$\frac{n^*}{L} = G(w^* - \min\{p, w\} - c) \quad (2)$$

Having the previous in mind, there is a unique pair (w^*, n^*) , that satisfies equations 1 and 2, which characterize the strategic equilibrium. Proposition 3 summarizes these findings.

Proposition 3. *There is a unique pair $(w^*, n^*) \in \mathbb{R} \times (0, L)$ that satisfies the system of equations 1 and 2, where a mass of n^* workers unionize, and the employer E and the labor union U reach an agreement w^* regarding unionized workers' wage without a strike. This pair characterizes the subgame perfect equilibrium of the game.*

2.2 The role of AI

We now study the effect of using adopting more AI on the bargaining power of labor unions. To do this we study changes to the price parameter p without loss of generality—thus the AI market is modeled in reduced form. First, there is a direct effect: AI affects how costly a strike is for the employer. Indeed, an increase in AI which is captured by an innovation that reduces prices ($\Delta p < 0$) reduces the share of the output that the union can secure for its members via wages, because by the law of demand $\partial r / \partial p < 0$.⁴

Second, there is an (strategic) indirect effect: AI substitutes workers thus strikes are less costly to employers weakening the bargaining power of labor unions, pushing down unionized workers' equilibrium wage and thus reducing the benefits of joining the union. As a result, the mass of workers that unionize, falls.

³Since each worker has a mass of 0, whether they join the union or not does not affect the union's bargaining power; to study collective action problems in union membership (such as free-riding) is outside of the scope of this paper.

⁴Assuming that there is no strategic effect (i.e. $\partial n / \partial r = 0$), then

$$\frac{\partial}{\partial r} [F(r+L) - F(r+L-n)] = F'(r+L) - F'(r+L-n) < 0.$$

Proposition 4 summarizes these findings.⁵

Proposition 4. *Assume (w^*, n^*) satisfy the equilibrium conditions (1) and (2), and assume that the pure economic benefit of belonging to the union is higher than that of not joining the union, $w^* - c > \min\{p, w\}$, then*

$$\frac{\partial n^*}{\partial r} \cdot \frac{\partial r}{\partial p}, \frac{\partial w^*}{\partial r} \cdot \frac{\partial r}{\partial p} > 0.$$

2.2.1 Endogenous demand for labor

Consider that before workers unionize, the employer chooses the mass of workers (L) knowing that some of them might, indeed, unionize. The employer's problem is therefore:

$$\begin{aligned} \max_{L \geq 0} & F(r + L - n^*(L)) - w(L - n^*(L)) - w^*(L)n^*(L) - r \cdot p \\ \text{s.t. } & (w^*(L), n^*(L)) \text{ satisfy (1) - (2).} \end{aligned}$$

In this case there are two channels through which AI affects unionization: i) The direct effect of AI on bargaining power, and ii) The employer's demand for labor.

Denoting $L^*(p)$ as the optimal demand for labor, the overall effect of AI on unionization is given by:

$$\begin{aligned} \frac{\partial}{\partial p} [n^*(L^*(r), r)] &= \overbrace{\frac{\partial n^*}{\partial L} \frac{\partial L^*}{\partial r} \frac{\partial r}{\partial p}}^{\text{effect through employment}} + \underbrace{\frac{\partial n^*}{\partial r} \frac{\partial r}{\partial p}}_{\text{effect through bargaining power}}. \end{aligned}$$

We know that the second term on the right-hand side of the equation is positive by Proposition 4, furthermore $\partial L^* / \partial r < 0$ because AI substitutes labor. To sign the effect of this derivative we assume that more workers entails more unionized workers—but this relationship is not one-to-one. Furthermore, $\lambda \rightarrow 1$ for simplicity.

⁵The assumption that the economic benefits of joining a union are greater for unionized workers is not necessary for the comparative statics, but it is a sufficient condition. In this regard, there is strong empirical evidence in support of this assumption (Farber et al., 2018).

Intuitively, we are assuming that some newly hired workers have incentives to join the labor union ($\partial n^*/\partial L > 0$). Moreover, both the labor union and the employer become more patient as their discount factor of the future converges to 1, but this happens in such a way that labor unions become more patient than the employer.⁶ Proposition 5 summarizes our findings.

Proposition 5. *Assume $w^* - c > \min\{p, w\}$, that $\partial n^*/\partial L > 0$, and that $\lambda \rightarrow 1$, then*

$$\frac{\partial}{\partial p} [n^*(L^*(r), r)] = \frac{\partial n^*}{\partial L} \frac{\partial L^*}{\partial r} \frac{\partial r}{\partial p} + \frac{\partial n^*}{\partial r} \frac{\partial r}{\partial p} > 0.$$

2.2.2 Effect of AI on union density n/L

It is standard to study unionization using the ratio of unionized workers to the total number of workers. This measure is also theoretically appealing because AI can affect both the numerator n^* and the denominator L , and both can decrease as the number of AI increases under certain (arguably commonplace) conditions (Proposition 5). Note, however that

$$\frac{\partial}{\partial r} \left[\ln \left(\frac{n^*(L^*(r), r)}{L^*(r)} \right) \right] = \frac{\frac{\partial n^*}{\partial L} \frac{\partial L^*}{\partial r} \frac{\partial r}{\partial p} + \frac{\partial n^*}{\partial r} \frac{\partial r}{\partial p}}{n^*} - \frac{\frac{\partial L^*}{\partial r} \frac{\partial r}{\partial p}}{L},$$

thus the sufficient condition that signs this derivative requires that incoming workers unionize at a higher rate than those who were already employed (Proposition 6). Therefore union density will decline under much more stringent conditions vis-à-vis union membership for instance.

Proposition 6. *If $\frac{\partial n^*}{\partial L} \geq \frac{n^*}{L^*}$, then $\frac{\partial}{\partial r} \left(\frac{n^*}{L^*} \right) \frac{\partial r}{\partial p} > 0$.*

2.2.3 Heterogeneous labor

We adopt the stylized fact that workers with higher marginal productivity receive higher non-union wages w vis-à-vis those with lower marginal productivity. In other words, the cost of opportunity for high productivity workers of joining a union is higher in comparison to lower productivity ones.

Proposition 7 establishes the relevant comparative statics with respect to w : It states that if the wages of non-unionized workers increase, the number of unionized workers decreases. This occurs because an increase in the benefits of the outside option (not unionizing) increases the opportunity cost of joining the union, pushing a fraction of unionized workers to de-unionize. Furthermore, the marginal effect of one more robot on the number of unionized workers intensifies as w increases.

⁶To assume $\lambda \rightarrow 1$ is equivalent to assuming that $\delta_i, \delta_E \rightarrow 1$, and $\frac{1-\delta_i}{1-\delta_E} \rightarrow 0$.

Also, as automation enter the work place, more productive workers are more likely to opt-out of their union because the opportunity cost of unionizing goes up.⁷

Proposition 7.

$$\frac{\partial n^*}{\partial w} < 0.$$

Moreover, if G is unimodal, then

$$\frac{\partial}{\partial w} \left(\frac{\partial n^*}{\partial r} \frac{\partial r}{\partial p} \right) > 0$$

if and only if $(\lambda F'(r + L - n^*) + (1 - \lambda)c - w^*)g(w^* - \min\{p, w\} - c) > \frac{n^*}{L}$.

2.3 Illustration

We illustrate the implications of the results above in Figure 1, below. In it we plot unionization as a function of the price of AI (solid line), and as a function of the the market wage (dashed line). These curves reflect propositions 6-7, which indicate that unionization increases with the price of AI and that it decreases with the market wage. This also implies that collective bargaining increases with the price of AI and it decreases with the market wage because this affects unions' bargaining power similarly. In other words, workers lose bargaining power when automation is cheaper and when the opportunity cost of unionizing is high, both of which translate into lower unionization and thus lower collective bargaining.

If there are no market frictions the curves intersect at $p = w$ because—under the shadow of strikes—firms wouldn't pay more than the market wage of a worker for a robot (and vice-versa), because AI perform the same tasks as a worker and because machines don't threat with strikes. Hence first, if automating is cheap at p' , the demand for them will go up because firms can use cheaper machines to replace workers. However, as the demand for AI goes up, it also increases its price until the indifference condition $p = w$. Note that this also implies that as the price adjustment takes place, workers opportunity cost from unionization falls and as such they “preemptively” start unionizing, going from n to n^* . Second, unionization at n'' is sub-optimal because under the shadow of strikes, unionized workers have bargaining power that allows them to obtain w^* .


⁷The unimodality assumption on G should not be restrictive in this case insofar as it is reasonable to assume that the modal worker cares more about whether they obtain material benefits from unionizing in comparison to the idiosyncratic factors that entail joining the union.

These results imply that when there is an innovation that facilitates AI, workers will start unionizing and bargaining preemptively before the innovation is implemented.

Note also that replacing productive workers at market wage w' is more difficult, because the price of machines needs to go up until $p = w'$. In other words, more AI adoption is needed to perform the tasks by a productive worker. This implies in turn that the opportunity cost of unionization for these workers is higher because it is harder for firms to replace them with machines, which is captured by the smaller slope of the curve $n(p)'$ versus the $n(p)$ curve. Therefore we should expect a lower collective bargaining from unions when there are more productive workers relative to less productive ones. A similar logic holds if there are labor market frictions such that w' is lower than w .

All in all, it also follows from the propositions above that if the opportunity cost of unionization increases—which can happen with higher values of ε_i , c , and w —then collective action from labor decreases. This implies in turn that worker heterogeneity in these parameters may also increase the opportunity cost of bargaining for workers.⁸

Figure 1: Changes in bargaining power as a result of changes in AI



⁸Balcazar (2023) and Balcazar (2025) shows that worker heterogeneity is key to understand the impact of AI on politics. Bisbee and Rosendorff (2024) show that labor market frictions can moderate the impact of economic shocks on political preferences.

3 Data and variables

The use of commercial and scalable AI is a relatively new phenomenon. The history of AI indicates that there was relatively low interest in AI until 1993, except for academic- and government-led interest that started in the late 50s. Later on, during the 2000s AI developments were mostly experimental—like Deep Blue and AlphaGo, which pitted the best chess and go players against machines. Other developments on pattern recognition, speech recognition, and spatial recognition became commercially scalable for providing services such as Alexa, Siri and other, and appliances such as vacuum cleaning robots like Roomba. However AI lacked the level of development to generate the impact we have seen since the early 2020s ([Achiam et al., 2023](#)).


The possibility of using AI in firms and other forms of production became more feasible when limitations in computer power, especially GPUs, were overcome, allowing the use of deep neural networks, and later Convolutional Neural Networks and Artificial Neural Networks during the early 2010s. However, it was not until the creation of the transformer architecture, which debuted in 2017, that tech-companies produced impressive generative AI applications using large language models (LLMs). The transformer architecture converted text into numerical representations using embeddings, via a parallel multi-head attention mechanism, leading to the development of scalable AI tools such as ChatGPT, Gemini, Claude, DeepSeek and other—some of which have been tailored for companies’ specific needs.

While AI’s development took place over decades, the release of ChatGPT was the milestone that catapulted AI into the mainstream as we see it today. To show this, I use normalized Google trends as a measure of interest for keywords such as GPT, ChatGPT and AI, noting a spike in interest at the release of ChatGPT, and later on around the release of GPT4 on March 14 of 2023—Figure 2, panel a. (I describe this measure in Appendix A.) Similarly, Figure 2, panel b, shows the news interest regarding the keywords above using the Global Database of Events, Language and Tone (GDELT), which is a real-time database of newsworthy events. It shows that news coverage about GPTs paints a similar picture, despite we observe a lag vis-à-vis Google trends, potentially due to idiosyncrasies about news-market dynamics that are outside of the scope of this paper.

I choose the release of ChatGPT 3.5. (ChatGPT from hereon) on November 30th as the event because we observe evidence of anticipation for the latter releases such as GPT4 in Figure 2. This implies sorting, which violates the assumptions for carrying out event studies, making the latter event unfit for analysis ([Muñoz et al., 2020](#)). Furthermore, too little time has taken place from ChatGPT’s release until today to do a proper analysis of its effects in comparison to studies on the effects of robot adoption, which use data that can span one or two decades, or even longer.⁹

⁹Long-term data is often necessary to disentangle the short- from the long-run effect of the shift using shift-share

Figure 2: Normalized Google trends interest scores for ChatGPT and AI



Note: Google trends interest scores for each country in the world are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. GDELT uses data from the Internet Archive’s Television News Archive. I focus on data from networks with international coverage such as CNN, Al Jazeera, BBC news, DeutscheWelle, Russia Today, ABC, CBS, NBC, Telemundo and Univision. Gray shaded areas are for 30-days after the release of ChatGPT and GPT4 respectively.

After checking data from various countries in North America, Latin America, Africa, Europe and Asia, I find that the data from the US is rich enough for properly analyzing this event and carrying out a proper statistical analysis, given availability of high-frequency data collected about the event that also provides high statistical power. Further, we also observe higher levels of inquiry about how to join a labor union around the advent of ChatGPT (Figure 3), indicating that the theorized mechanism can be at play. Hence from here onwards I restrict myself to US data.

3.1 Collective Bargaining Agreements (CBAs)

In the US, employers and labor unions are required by law to collectively bargain a contract when there are disagreements; they must also agree on the duration of such contract. Collective bargaining agreements intend to provide workers (or employers) with opportunities to negotiate salaries, staffing decisions, production decisions, etc. When a contract under collective bargaining agreement expires, or one of the parties breaks the contract, the parts have to collectively bargain another contract. If the the parts involved do not reach an agreement, workers can go on strike. This is commensurate with the bargaining model above.

The data on collective bargaining comes from the U.S. Federal Mediation and Conciliation Service designs (Jaeger et al., 2018).

Figure 3: Normalized Google trends interest scores for unionization and AI threat



Note: Google trends interest scores for each the US are obtained from using ChatGPT as a keyword since data becomes available. This is normalized by term where the highest interest = 100. Gray shaded areas are for 120-days after the release of ChatGPT and GPT4 respectively.

vice (FMCS). These data has information about ongoing process at the date of the event, and new processes after the event. It also contains information about the geographical location and industry of the collective bargaining process, and the number of workers in the collective bargaining unit.

The National Labor Relations Act requires that employers or the representative of the employees who are covered by a collective bargaining agreement provide written notice to the Federal Mediation and Conciliation Service (FMCS) and a similar state agency—for instance in New Jersey it would be the State Board of Mediation (NJSBM). Failure to comply can result in legal and financial consequences. Hence the data arguably includes the universe of collective bargaining notices.

Standard data on unionization comes from yearly reports and it is not available at a high frequency to properly study the impact of ChatGPT using an event study design, due to potential anticipation effects regarding subsequent releases. These data may also induce lack of comparability about the event bandwidth. In particular, data from the Current Population Survey—which is used to measure unionization in the US—does not provide data about the day of the interview, but only the month which could induce violation of the ignorability assumption in event studies and make the analysis of manipulation unreliable. Nonetheless, I use this data to evaluate the degree

of anticipation via unionization.

I investigate the impact of ChatGPT on new and ongoing collective bargaining agreements at the zip code level. This also allows me to exploit any potential anticipation as a feature of the empirical design. I measure the number of bargaining agreements, the size of the bargaining unity, and the number of days from initiation to expiration.¹⁰ Table A1, in the Appendix, displays the summary statistics for all variables restricting the data to after the release of ChatGPT.

3.2 Strikes

Data on work stoppages is unavailable. In 2020, the FMCS stopped publishing similar data on strikes and removed the archival data from their website.¹¹ Data on strikes gleaned from the Bureau of Labor statistics, unfortunately, don't have good geographical information for the analysis herein. This precludes me from analyzing this outcome.

3.3 ChatGPT

I use Google searches as a proxy for exposure to the AI innovation because it reflects the extent of people's interest (excitement, worry, etc.) in regards to ChatGPT, which is a strategy that has been used to measure exposure to other events (Mellon, 2013). Using news supply doesn't reflect this choice insofar as we observe a substantive lag in news coverage (Figure 2), instead reflecting news market dynamics. Google trends reflect a costly choice resulting from the event, to become informed about it. Therefore it is a proxy of the average level of treatment uptake—I provide further evidence below.

Google trends also allow us to obtain subnational level data at the level of the Designated Market Area (DMA), which facilitate a granular analysis.¹²

Figure 4, panel a, gives credence to the quasi-exogenous nature of the event in the US' context, as we observe a spike in interest at the release of ChatGPT. Panel b shows that lay-offs increased about the advent of ChatGPT—thus the shock was indeed economically relevant. Panel c displays the geographical distribution regarding interests in ChatGPT. We observe important geographical variation, with higher interest scores in California, Texas and the East Coast where the tech industry concentrates (panel c). This, however, isn't correlated with subsequent job vacancies in AI-related


¹⁰The share of workers in the bargaining unit cannot be measured because the data on establishment size stopped being collected in September of 2022.

¹¹FOIA requests to obtain more recent data have been denied.

¹²There are 210 media markets as defined by Nielsen.

jobs (Figure 4, panel c),¹³ indicating that the statistical compliers may not necessarily be those most impacted by AI itself. This echoes recent evidence showing that the risk to be replaced by AI may have little explanatory power on perceptions of AI threat (Menon and Zhang, 2024; Anelli et al., 2025; Baccini et al., 2025).

**Figure 4: Interest scores for ChatGPT
(United States)**



Note: Google trends interest scores for each DMA in the US are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. This is normalized by the top searched terms for comparability. For private sector lay-offs I use data from JOLTs. For AI job vacancies I use data from Lightcast (Acemoglu et al., 2022).

¹³For AI job vacancies I use data from Lightcast—a leading source of labor market information with a vast repository of millions of job postings. Further details on this data can be found in Acemoglu et al. (2022).

3.3.1 Measuring the threat from AI

The literature about AI often assumes that AI creates anxiety among workers fearing being replaced and thus it will affect politics (e.g., [Borwein et al. 2024](#); [Magistro et al. 2024](#)). However, [Chaudoin and Mangini \(2022\)](#) shows that this may depend on the content of elite cues about automation. In this regard, recent analyses of AI narratives indicate that they have not necessarily been negative or related to a higher risk of unemployment. Instead, they tend to be generally positive and related to the benefits of technological development ([Fast and Horvitz, 2017](#); [Neri and Cozman, 2020](#); [Kelley et al., 2021](#); [Bareis and Katzenbach, 2022](#); [Reauscher, 2022](#); [Roe and Perkins, 2023](#); [Ryazanov et al., 2025](#); [Wu, 2024](#)); when AI narratives are negative, they are associated to a potential existential threat where AI can deem humanity unnecessary.

Recent observational evidence supports this argument. [Antoniades et al. \(2025\)](#), for instance, finds that AI can benefit the incumbent because it has come accompanied by higher employment, especially for highly educated individuals—see also [Antoniades et al. \(2024a,b\)](#). There is no evidence that AI has generated widespread unemployment (e.g., [Aghion and Bunel 2024](#); [Acemoglu 2025](#); [Brynjolfsson et al. 2025](#)).

To address this conundrum I borrow from the literature on “availability bias” ([Tversky and Kahneman, 1973, 1974](#)), to define the extent of exposure to AI.

The heuristic of availability bias implies that when faced with new or familiar situations—like losing a job, getting ill, facing a natural disaster, etc.—people assign higher probabilities to outcomes coming from salient events that share similar features. It is a cognitive shortcut we use naturally when making decisions under uncertainty ([Griffin et al., 2001](#); [Weyman and Barnett, 2016](#)). The ease with which consequences can be visualized, particularly if they are widespread, catastrophic or enduring, can elicit a myopic response that inflates probability estimates of outcomes coming from events that are believed to share one or more features. These effects have been widely cited as sponsoring “read-across” from one risk agenda to another ([Marx et al., 2007](#)).¹⁴

Although read-across effects can be an inappropriate representation of the event itself, they are indicative of the features of the issue or event that are considered to be most salient for the person. Thus whether the event is being interpreted correctly is a moot point. If one’s job was threatened, or disappeared as a result of automation, one may use this experience to interpret a similar event like the advent of a new technology. These heuristics also have spillovers in groups with smaller stock of cognitive resources ([Altmann et al., 2022](#)). Thus the heuristic may be visible in groups sharing similar characteristics—like groups of workers—even if some members didn’t experience

¹⁴People may also underestimate these probabilities if recalling analogous events is limited, or when no personal experience or historical precedents exist for the event. In this sense, an unlikely event is one that is hard to imagine.

the events directly.

In the context herein, workers in industries that have been previously exposed to ICTs should be more likely to have been exposed directly (or indirectly through socialization at the workplace) to the prospect of technological unemployment as a result of new computer technologies. I assume that the extent to which the heuristic is available to an individual can be captured by earlier measures of exposure to ICTs such as the Routine Task Intensity (RTI) measure by Autor (2013), which measures how computer technology altered the demand for job skills. This measure exhibits low correlations with similar task routine measures and those that are specific to AI (Appendix B).

Treatment uptake. I construct a sector-level measure of task routinization using population cells to increase variation and reduce measurement error. I use age, education, industry, gender, ethnicity and county to create the cells. The measurement assumption is that cell averages are representative of individuals with similar characteristics, which I guarantee by using US census data for 2010 to construct the weights, to then obtain weighted cell averages.

Then I construct and analyze profiles for these indexes. I show that young individuals are more likely to be exposed to ICTs, and those with higher education, less. This stands in contrast to AI-specific indexes, where educated individuals are more likely to be exposed to AI.

Validation. The identifying assumption for the mechanism herein is that, Google searches capture interest, worry, etc., of workers exposed to these technologies. This is warranted insofar as the advent of ChatGPT coincided with increased lay-offs in the private sector as I show above.

To validate the assumption I use data from the 119 wave of the American Trends Survey by the Pew Research Center, which was conducted from Dec. 12 to Dec. 18, 2022. The data contains numerous questions about respondent's knowledge and concern regarding AI technologies. I map both the Google scholar searches and the RTI measures cells described above to this data set for analysis.

I present the descriptive analysis in Appendix C. First, I find that respondents that have accurate knowledge about AI report that they have used it, or that they have heard about it; they are also more likely to use Internet and social media. This implies, as per the previous findings, that young individuals are both more likely to be exposed to ICTs and know more about AI.¹⁵

Second, individuals living in areas with high Google search scores for ChatGPT and higher levels of RTI are more likely to know about AIs.¹⁶ This indicates that Google search scores likely reflect exposure to and selection into the treatment uptake. Therefore the combination of Google

¹⁵Unfortunately, the survey doesn't have data on union status to check if union members have higher or lower knowledge about and exposure to AI than non-union ones.

¹⁶This may explain the lack of correlation between Google searches and AI job vacancies in Figure 4, panel d.

searches and RTI should allow us to focus the analysis on the statistical compliers (e.g., [Anelli et al. 2021](#)).

4 Empirical design

To estimate the effect of ChatGPT on the outcomes I use an event-study design. The advantage of this approach is that we can control for location and period fixed effects, which can address concerns about seasonality or time-invariant confounders.

I estimate the following regression equation:

$$Y_{iudtp} = \beta \text{ChatGPT}_{idt} + X'_{iudtp} \delta + f(t; r)' \gamma + \theta_d + \omega_u + \rho_p + \varepsilon_{iudtp}, \quad (3)$$

where Y_{iudtp} is the outcome for zip code i , in DMA d , 2-digit industry u , date t , and year p ; ChatGPT_{idt} is the normalized level of interest in ChatGPT; $\mathbf{X}_{dtp}$ is a vector of pre-treatment night-light measures to proxy for changes in economic activity,¹⁷ Covid cases from the CDC¹⁸, and public sector status; μ_d , ω_u and ρ_p are DMA, industry and survey-round fixed effects; $f(\cdot)$ is some flexible function of time about the release date r which also captures seasonality; and ε_{iudtp} denotes the idiosyncratic error term. Importantly, I define t at the week-level as I observe substantial week-level seasonality. I use doubly-robust estimation to address the issue of negative weights ([Arkhangelsky et al., 2024](#)). Standard errors are spatially clustered.

The event-study plot (Figure 5) gives credence of strong albeit transitory effect.


To estimate the regression equation above, I first estimate an optimal data-driven bandwidth around the release of ChatGPT ([Cattaneo et al., 2019, 2024](#); [Calonico et al., 2025](#)). This analysis is presented in Appendix D, where I report both mean squared error (MSE) and coverage error-rate (CER) optimal bandwidths, both one and two sided. Given the results, I choose a 11-week bandwidth. I also carry out balance, anticipation and manipulation checks, indicating that the data before and after the event are statistically comparable within the bandwidth. I also explore potential anticipation, finding no evidence for it.

One of the operational advantages of this design is that many measurable/observable economic aggregates that could confound the effect of the event are held constant. On the one hand, this is circumstantial insofar as most economic aggregates are measured at the quarterly or yearly level,

¹⁷I use version 2 VIIRS nighttime lights (V2 VNL) ([Gibson and Boe-Gibson, 2021](#)).

¹⁸https://data.cdc.gov/Public-Health-Surveillance/United-States-COVID-19-Community-Levels-by-County/3nnm-4jni/about_data

Figure 5: Event study plot for the effect of ChatGPT release on collective bargaining agreements



Note: 95% Conley confidence intervals in parentheses. Estimates are doubly robust. Controls include: nightlights as a proxy of economic activity, the number of covid cases, public sector status, the China import shock, the robot adoption shock, and offshoring. Regressions use observations in a 30-day optimal bandwidth about the release of ChatGPT.

and rarely at a lower level of time-disaggregation, much less so at the subnational level. On the other hand, these confounders vary little month-to-month. Nevertheless, I proxy monthly variation in economic activity with nightlights.¹⁹

β is the coefficient of interest, and it measures the effect of one standard deviation in $ChatGPT_{idt}$. Since ‘ChatGPT’ is searched 30.4 million times per month, with 3 to 4 searches on average per person, a standard deviation is equivalent to $30.4/(4 \times 30) = 0.256$ million individuals searching the term per day using Google.²⁰ I anticipate $\beta > 0$.

5 Results

I first analyze the average impact of the advent of ChatGPT on the collective bargaining outcomes. Table 1 shows the results. In Column (1) we observe that an increase in one standard deviation in the ChatGPT shock increases the number of collective bargaining agreements by 26.4 percentage points (PPs); this is about one tenth of a standard deviation in the outcome. In Column (3) we find that the shock increases the likelihood of starting a new collective bargaining process by 3.1 PPs. We do not observe statistically significant effects on the size of the bargaining unit or the number

¹⁹The GDP-luminosity elasticity is usually between 0.1 and 0.25.

²⁰<https://www.semrush.com/blog/most-searched-keywords-google/>

of days from initiation to termination of collective bargaining agreements.

In Appendix D.1 I evaluate the stability of parameter's to anyone DMA using Jackknife, finding that the results are robust.

Table 1: Effect of ChatGPT release on collective bargaining agreements

	Log (outcome +1)			
	<i>Number of collective bargaining cases</i>	<i>Bargaining unit size</i>	<i>Number of days from initiation to expiration</i>	<i>New bargaining process</i>
	(1)	(2)	(3)	(4)
ChatGPT	0.2639** (0.1063)	0.1552 (0.1455)	0.0684 (0.0617)	0.0308* (0.0177)
Observations	4312	4251	3980	4312
County fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes

Notes: Conley standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Estimates are doubly robust. Controls include: nightlights as a proxy of economic activity, the number of covid cases, and public sector status. Regressions use observations in a 30-day optimal bandwidth about the release of ChatGPT.

Albeit we shouldn't expect a similar reaction from public sector unions, the percentage of collective bargaining agreements is 3% which doesn't lend itself to a meaningful statistical analysis due to the small number of observations. The results presented herein, however, exclude these type of unions.

Heterogeneous effects. Next, I investigate if the pre-treatment level of exposure to ICTs moderates the effect above. This allows me to evaluate the role of worker heterogeneity using the availability heuristic, providing a measure of treatment uptake. I estimate the following extended regression equation:

$$Y_{iudtp} = \beta_1 \text{ChatGPT}_{idt} + \beta_2 \text{ChatGPT}_{idt} \times \text{RTI}_{iudt} + \beta_3 \text{ChatGPT}_{idt} \times Z_{iudt} + X'_{iudtp} \delta + f(t; r)' \gamma + \theta_d + \omega_u + \rho_p + \varepsilon_{iudtp}, \quad (4)$$

where Z_{iudt} is a vector of pre-treatment confounding mechanisms such as robot adoption, offshoring and import competition (the China shock). The estimation procedure is the same as above.

The coefficient of interest is $\beta_1 + \beta_2 \times \text{RTI}_{iudt}$. β_1 is the effect of an increase in one standard deviation in the ChatGPT-level of exposure when no task can be routinized by ICTs: $\text{RTI}_{iudt} = 0$. $\beta_1 + \beta_2 \times \text{RTI}_{iudt}$ is the effect of an increase in one standard deviation the ChatGPT-level of exposure when RTI_{iudt} takes some value in its support. I anticipate $\beta_2 > 0$.

Figure 6, Panel A, shows that the number of collective agreements are usually higher in response to the advent of ChatGPT for higher pre-treatment levels of exposure to ICTs. In panels B and C we also observe that bargaining unit size and the number of days from initiation of the collective bargaining agreement, to its expiration, increase for the highest quartile of the distribution of pre-treatment levels of exposure to ICTs, in response to the shock. However, we do not observe a heterogeneous effect for the likelihood of new collective bargaining agreements (panel D).

All in all, these results give credence to the idea that the observed results are driven by workers in industries that were previously exposed to ICTs.


6 Conclusions

The advent of generative AI has had a profound effect on society, but still little is know about its effect on collective action besides anecdotal evidence on worker strikes regarding AI implementation. Studying the effect of the advent of ChatGPT using the US as a case study, I demonstrate that this technological innovation increased the level of collective action from labor through both higher levels of collective bargaining agreements between firms and workers. I establish a theory of collective bargaining that indicates that this occurs because workers preemptively mobilize while their bargaining power is higher, because once AI is implemented through increased capital deepening their bargaining power decreases.

I also show that workers in industries previously exposed to changes in demand for skills due to computer technologies, are more likely to respond to the ChatGPT innovation. Therefore I show that studying worker heterogeneity is essential to understand the impacts of AI on behavior.

As AI continues to develop, changing society, economics and politics, the findings herein demonstrate that technological innovation shapes collective action. This echoes previous work on the impact of robot adoption on union power (Balcazar, 2023; Becher and Stegmueller, 2025; Anelli et al., 2025). Specifically, I demonstrate that while AI may have the potential to replace the tasks performed by millions of workers, it also changes the strategic incentives in the distributive conflict between workers and firms, leading to higher levels of union activity through collective bargaining, even if workers' bargaining power can be negatively affected by automation.

Figure 6: Heterogeneous effect of ChatGPT release on collective bargaining agreements



Note: 95% Conley confidence intervals in parentheses. Estimates are doubly robust. Controls include: nightlights as a proxy of economic activity, the number of covid cases, public sector status, the China import shock, the robot adoption shock, and offshoring. Regressions use observations in a 30-day optimal bandwidth about the release of ChatGPT.

All in all, the paper challenges the notion that union activity has been dwindling as automation has incentivized worker mobilization inside the work place. It also challenges the notion that marked fundamentalism is becoming the fundamental ideology when workers assess the relationship with the firm in a changing economic landscape. It shows that technological change affects workers incentives to overcome collective action issues to pursue their interests within the work place. Although exploring the effect of AI on political mobilization from workers is outside of the scope of this research, the findings herein indicate that this is an agenda worth pursuing.

References

Acemoglu, D. (2025). The simple macroeconomics of ai. *Economic Policy*, 40(121):13–58.

- Acemoglu, D., Aghion, P., and Violante, G. L. (2001). Deunionization, technical change and inequality. In *Carnegie-Rochester conference series on public policy*, volume 55, pages 229–264. Elsevier.
- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Achiam, J., Adler, S., Agarwal, S., et al. (2023). Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Aghion, P. and Bunel, S. (2024). Ai and growth: where do we stand? *Policy Note*.
- Agnolin, P., Anelli, M., Colantone, I., and Stanig, P. (2025). Robots replacing trade unions: Novel data and evidence from western europe.
- Alhquist, J. and Downey, M. (2021). The effects of import competition on unionization. UCSD mimeo.
- Altmann, S., Grunewald, A., and Radbruch, J. (2022). Interventions and cognitive spillovers. *The Review of Economic Studies*, 89(5):2293–2328.
- Anelli, M., Colantone, I., Gallego, A., and Stanig, P. (2025). What if you see it?workers’ perceptions of and reactions to llms. *Presented at AIPSR in Barcelona*.
- Anelli, M., Colantone, I., and Stanig, P. (2018). We were the robots: Automation and voting behavior in western europe. Technical Report 2019-115, Bocconi University–Mimeo.
- Anelli, M., Colantone, I., and Stanig, P. (2021). Individual vulnerability to industrial robot adoption increases support for the radical right. *Proceedings of the National Academy of Sciences*, 118(47):e2111611118.
- Antoniades, A., Balcazar, C. F., Chatzikonstantinou, E., and Kern, A. (2025). Did the advent of ai offset the far-right? evidence from us’ elections.
- Antoniades, A., Chatzikonstantinou, M., and Ahman, M. S. (2024a). The convergence of occupations: Evidence from online job posts.
- Antoniades, A., Chatzikonstantinou, M., and Savka, O. (2024b). Are a.i. jobs driving up the demand for stem education? yes, but not for women.

- Arkhangelsky, D., Imbens, G. W., Lei, L., and Luo, X. (2024). Design-robust two-way-fixed-effects regression for panel data. *Quantitative Economics*, 15(4):999–1034.
- Autor, D. (2013). The “task approach” to labor markets: an overview. *Journal for Labour Market Research*, 46(3):185–199.
- Baccini, L., Kleinberg, K., and Weymouth, S. (2025). Technology, labor shocks, and support for liberalmarket institutions. *Presented at AIPSR in Barcelona*.
- Balcazar, C. F. (2023). Globalization, unions and robots: The effects of automation on the power of labor and policymaking. *Available at SSRN 4574527*.
- Balcazar, C. F. (2025). A theory of distributive conflict for the age of ai. *Mimeo, University College London*.
- Balcazar., C. F. (2025). A theory of distributive conflict for the age of ai. *Mimeo, University College London*.
- Bareis, J. and Katzenbach, C. (2022). Talking ai into being: The narratives and imaginaries of national ai strategies and their performative politics. *Science, Technology, & Human Values*, 47(5):855–881.
- Becher, M. and Stegmueller, D. (2021). Reducing unequal representation: The impact of labor unions on legislative responsiveness in the u.s. congress. *Perspectives on Politics*, 19:92–109.
- Becher, M. and Stegmueller, D. (2025). Machines againstworkers? the heterogeneous impact ofrobots on union strength. *Mimeo, American Political Science Association*.
- Becher, M. and Stegmuller, D. (2023). Global competition, local unions, and political representation: Disentangling mechanisms. Technical report, Unpublished manuscript.
- Bisbee, J. and Rosendorff, B. P. (2024). Antiglobalization sentiment: Exposure and immobility. *American Journal of Political Science*.
- Borwein, S., Bonikowski, B., Loewen, P., Magistro, B., and Lee-Whiting, B. (2024). Who can assert ownership over automation? workplace technological change, populist and ethno-nationalist rhetoric, and candidate support. *Political Behavior*.
- Brynjolfsson, E., Chandar, B., and Chen, R. (2025). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Technical report, Working paper. Latest version available at <https://digitaleconomy.stanford> . . .

- Calonico, S., Cattaneo, M. D., Farrell, M. H., Palomba, F., and Titiunik, R. (2025). Treatment effect heterogeneity in regression discontinuity designs. *arXiv preprint arXiv:2503.13696*.
- Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2019). *A practical introduction to regression discontinuity designs: Foundations*. Cambridge University Press.
- Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2024). *A practical introduction to regression discontinuity designs: Extensions*. Cambridge University Press.
- Chatterjee, S. (2021). A new coefficient of correlation. *Journal of the American Statistical Association*, 116(536):2009–2022.
- Chaudoin, S. and Mangini, M.-D. (2022). Robots, foreigners, and foreign robots: Policy responses to automation and trade. Mimeo, Harvard University.
- Dean, A. (2025). If there is no struggle, there is no progress: Labor-led democratization and labor power. Mimeo George Washington University.
- Eggers, A. C., Freier, R., Grembi, V., and Nannicini, T. (2018). Regression discontinuity designs based on population thresholds: Pitfalls and solutions. *American Journal of Political Science*, 62(1):210–229.
- Farber, H. S., Herbst, D., Kuziemko, I., and Naidu, S. (2018). Unions and inequality over the twentieth century: New evidence from survey data. Technical Report 3, National Bureau of Economic Research.
- Fast, E. and Horvitz, E. (2017). Long-term trends in the public perception of artificial intelligence. Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence.
- Frey, C. B. (2019). *The Technology Trap*. Princeton University Press.
- Frey, C. B. and Osborne, M. A. (2017). The future of employment: how susceptible are jobs to computerisation? *Technological forecasting and social change*, 114:254–280.
- Gibson, J. and Boe-Gibson, G. (2021). Nighttime lights and county-level economic activity in the united states: 2001 to 2019. *Remote Sensing*, 13(14):2741.
- Gil, A. (1986). Some labour implications of technological change in rail and air transport. *Int'l Lab. Rev.*, 125:1.
- Gingrich, J., Wu, N., and Zhang, B. (2025). Workplace governance and labor perceptions of technological risks and benefits. *Industrial and Corporate Change*, page dtaf059.

- Grembi, V., Nannicini, T., and Troiano, U. (2016). Do fiscal rules matter? *American Economic Journal: Applied Economics*, pages 1–30.
- Griffin, D., Gonzalez, R., and Varey, C. (2001). The heuristics and biases approach to judgment under uncertainty. *Blackwell handbook of social psychology: Intraindividual processes*, 1:207–235.
- Haapanala, H., Marx, I., and Parolin, Z. (2023). Robots and unions: The moderating effect of organized labour on technological unemployment. *Economic and Industrial Democracy*, 44(3):827–852.
- Imbens, G. W. (2003). Sensitivity to exogeneity assumptions in program evaluation. *American Economic Review*, 93(2):126–132.
- Jaeger, D. A., Ruist, J., and Stuhler, J. (2018). Shift-share instruments and the impact of immigration. Technical report, National Bureau of Economic Research.
- Kaplan, E. and Naidu, S. (2025). Between government and market: The political economics of labor unions. *Annual Review of Economics*, 17.
- Kelley, P. G., Yang, Y., Heldreth, C., Moessner, C., Sedley, A., Kramm, A., Newman, D. T., and Woodruff, A. (2021). Exciting, useful, worrying, futuristic: Public perception of artificial intelligence in 8 countries. In *AIES 2021 - Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 627–637. Association for Computing Machinery, Inc.
- Kennedy, D. et al. (1982). Labor and technology: Union response to changing environments.
- Magistro, B., Loewen, P., Bonikowski, B., Borwein, S., and Lee-Whiting, B. (2024). Attitudes toward automation and the demand for policies addressing job loss: the effects of information about trade-offs. *Political Science Research and Methods*.
- Marx, S. M., Weber, E. U., Orlove, B. S., Leiserowitz, A., Krantz, D. H., Roncoli, C., and Phillips, J. (2007). Communication and mental processes: Experiential and analytic processing of uncertain climate information. *Global Environmental Change*, 17(1):47–58.
- Mellon, J. (2013). Where and when can we use google trends to measure issue salience? *PS: Political Science & Politics*, 46(2):280–290.
- Menon, A. and Zhang, B. (2024). Messaging about different ai development timelines fails to change the public’s policy preferences.

- Menon, A. and Zhang, B. (2025). Future shock or future shrug? public responses to varied artificial intelligence development timelines.
- Muñoz, J., Falcó-Gimeno, A., and Hernández, E. (2020). Unexpected event during survey design: Promise and pitfalls for causal inference. *Political analysis*, 28(2):186–206.
- Neri, H. and Cozman, F. (2020). The role of experts in the public perception of risk of artificial intelligence. *AI and Society*, 35:663–673.
- Owen, E. and Park, S. (2025). Offshoring and the decline of labor unions: How globalization fosters technological change. *Mimeo, University of Pittsburgh*.
- Reauscher, N. (2022). *The Future of Work in the United States: Discourses on Automation and the Platform Economy*. Springer.
- Roe, J. and Perkins, M. (2023). ‘what they’re not telling you about chatgpt’: exploring the discourse of ai in uk news media headlines. *Humanities and Social Sciences Communications*, 10.
- Rothstein, S. A. (2022). *Recoding power: tactics for mobilizing tech workers*. Oxford University Press.
- Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, pages 97–109.
- Ryazanov, I., Öhman, C., and Björklund, J. (2025). How chatgpt changed the media’s narratives on ai: a semi-automated narrative analysis through frame semantics. *Minds and Machines*, 35(1):1–24.
- Thewissen, S. and Rueda, D. (2019). Automation and the welfare state: Technological change as a determinant of redistribution preferences. *Comparative Political Studies*, 52(2):171–208.
- Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, 5(2):207–232.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, 185(4157):1124–1131.
- Weyman, A. and Barnett, J. (2016). Heuristics and biases in decision making about risk. In *Routledge handbook of risk studies*, pages 131–139. Routledge.

- Wu, N. (2022). Misattributed blame? attitudes toward globalization in the age of automation. *Political Science Research and Methods*, 10(3):470–487.
- Wu, N. (2024). Foreign competition and technological attitudes in the united states *.
- Zhang, B. (2022). No rage against the machines: Threat of automation does not change policy preferences. In *AIES 2022 - Proceedings of the 2022 AAI/ACM Conference on AI, Ethics, and Society*, pages 856–866. Association for Computing Machinery, Inc.

ONLINE APPENDIX

Contents

1	Introduction	1
2	Theory	4
2.1	Bargaining wages and unionizing under the shadow of strikes	6
2.2	The role of AI	7
2.2.1	Endogenous demand for labor	8
2.2.2	Effect of AI on union density n/L	9
2.2.3	Heterogeneous labor	9
2.3	Illustration	10
3	Data and variables	12
3.1	Collective Bargaining Agreements (CBAs)	13
3.2	Strikes	15
3.3	ChatGPT	15
3.3.1	Measuring the threat from AI	17
4	Empirical design	19
5	Results	20
6	Conclusions	22
A	Measurement appendix	A2
B	Routine task indexes	B1

C	Validating the measure of exposure	C1
D	Event study design diagnostics	D1
D.1	Robustness	D1
E	Descriptives	E1

A Measurement appendix

Since Google measures interest in an event using a normalized score that sets the highest level of interest in a given place and period of time to 100, I re-normalize the level of interest in ChatGPT using the most searched terms in the world around November 30th to make the measure more comparable across geographical areas: The most searched terms were world cup, Google, weather and youtube, especially the first one because it was a major global event taking place from November 20th to December 18th. Despite this, I find that ChatGPT was the 6th rising term—i.e., the term with 6th highest growth in search volume—and in the 5th percentile for overall trend rankings in December of 2022, indicating that this keyword had one of the most significant growths in Google searches during the time period.

We observe growing interest in ChatGPT and AI, noting a spike in interest at the release of ChatGPT, an later on around the release of GPT4 on March 14 of 2023 across all world regions (Figure A1). This normalization also shows evidence of substantial geographical variation (Figure A2). In particular, using data from the World Bank’s World Development Indicators, we observe that more developed countries—with higher shares of educated individuals, and with higher research output—display higher levels of interest in ChatGPT at its onset (Figure A3).


These numbers are meaningful insofar as Google processes 8.5 billion searches daily, with 3-4 searches per person, suggesting that 2 billion people use Google world wide—more than 20% of the worlds population or about 35% of all internet users.^{A.1} Assuming that the most searched term on Google is “youtube” with 1.38 billion monthly searches,^{A.2} a back of the envelope calculation suggest that close to $1.35/8.5 \times 2 = 320$ million people may have searched the term ChatGPT during the time frame. This is consistent with data indicating that the volume of monthly searches about ChatGPT on Google is 0.62 billion.^{A.3}

^{A.1}See for example <https://searchengineland.com/google-5-trillion-searches-per-year-452928> and <https://www.semrush.com/blog/google-search-statistics/>

^{A.2}<https://www.semrush.com/blog/most-searched-keywords-google/>


^{A.3}*Ibid.*

**Figure A. 1: Normalized Google trends interest scores for ChatGPT and AI
(world regions)**




Note: Google trends interest scores for each DMA in the US are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. DMA's are defined by Nielsen's ratings. Gray shaded areas are for 30-days after the release of ChatGPT and GPT4 respectively

Figure A. 2: Normalized Google trends interest scores for ChatGPT and AI, 30-day average interest



Note: Google trends interest scores for each country in the world are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. This is normalized by the top searched terms for comparability.

Figure A. 3: Residualized correlations with 30-day average interest in ChatGPT



Note: Google trends interest scores for each country in the world are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. This is normalized by the top searched terms for comparability.

B Routine task indexes

There are many routine task indexes in existence for measuring the extent to which ICTs and AI could replace a task performed by a worker. Other RTI indexes measure off-shoring, which has allowed researchers to investigate remote work because they capture off-site, no face-to-face tasks. Given these characteristics of remote work, these tasks could also be related to those that could be automatable by AI. However, ICT-specific indexes should be commensurate with availability bias. In particular, the RTI index by Autor (2013) measures how computer technology altered job skill demands.

I start by comparing this index to other ICT-specific, AI-specific and off-shoring specific indexes. To do this I collect comprehensive list of thirty RTI indexes and their sources Table B2. I map all indexes to the Standard Occupational Classification (SOC) codes to make them comparable. AD RTI, however, is the main index of interest.


First, in Figure B1 I display the Pearson correlation table for the indexes considered. The top row is the AD RTI. We observe that this index is weakly correlated to the alternative computerization indexes. For instance, Frey and Osborne (2017) computerization index exhibits a relatively low association with the latter index. Second, the AD RTI is negatively correlated with most AI-specific indexes, except for small small positive correlations. All in all, these results indicate that AI indexes contrast the AD RTI. Third, the index of interest is also weakly correlated with off-shoring indexes.

Using the generalized coefficient of correlation to measure the potential existence of non-linear associations (Chatterjee, 2021), I find that there is little evidence of this. We also see that some AI measures are more correlated with off-shoring ones, indicating that they can be re-expressed as non-linear transformations of each other.

Next we map the measures using occupational codes to the US 2010 census. Table B1 displays the population-weighted averages for different indexes by age, gender, ethnicity, education and industry. We observe that for AD RTI young individuals are less more likely to be exposed to automation; individuals with higher education are less likely to be exposed; workers in manufacturing and professional, scientific, and technical occupations are more likely to be exposed too. Note the AD RTI provides a strong contrast with manual and abstract tasks, which are less automatable.

Figure B. 1: Correlation matrices between indexes

(a) Linear



(b) Generalized




Table B. 1: Profiles for RTI indexes

	Computerisation		
	AD RTI routine	AD RTI manual	AD RTI abstract
Male	-0.025* (0.011)	0.572*** (0.010)	0.165*** (0.010)
A. Age			
18-29	0.176*** (0.016)	0.007 (0.011)	-0.345*** (0.033)
30-45	0.006 (0.014)	0.090*** (0.012)	0.122*** (0.026)
45-64	0.024*** (0.006)	0.048*** (0.009)	0.155*** (0.018)
B. Ethnicity			
White non-Hispanic	-0.169*** (0.035)	-0.058*** (0.012)	0.263*** (0.026)
Black non-Hispanic	-0.150*** (0.038)	-0.084*** (0.009)	-0.088** (0.029)
Hispanic	0.194*** (0.044)	-0.273*** (0.029)	0.432*** (0.042)
C. Education			
Highs school	-0.024 (0.024)	-0.235*** (0.022)	0.452*** (0.032)
Some college	-0.093** (0.033)	-0.407*** (0.022)	0.787*** (0.032)
Technical	-0.167*** (0.035)	-0.441*** (0.021)	0.934*** (0.031)
College	-0.420*** (0.036)	-0.616*** (0.016)	1.932*** (0.041)
Graduate	-0.586*** (0.028)	-0.718*** (0.030)	2.993*** (0.052)

	Computerisation		
	AD RTI routine	AD RTI manual	AD RTI abstract
E. Industry			
Hospitality, service, arts, entertainment	-0.604*** (0.048)	-0.762*** (0.034)	-0.208** (0.072)
Health care and social assistance	0.143*** (0.042)	-0.639*** (0.030)	0.253*** (0.053)
Manufacturing, mining, construction, agriculture	1.243*** (0.057)	-0.379*** (0.054)	0.344*** (0.072)
Retail and trade	-0.125* (0.065)	-1.096*** (0.028)	0.452*** (0.066)
Education	-1.268*** (0.053)	-0.377*** (0.036)	0.470*** (0.051)
Banking, accounting, finance, real state	-0.205** (0.079)	-0.945*** (0.038)	0.291** (0.087)
Transportation	-0.721*** (0.055)	0.306*** (0.048)	-0.813*** (0.074)
Government	0.225* (0.105)	-0.663*** (0.035)	0.443*** (0.065)
Professional, scientific, technical	1.332*** (0.061)	-1.058*** (0.036)	1.667*** (0.049)
Observations	5030676	5030676	5030676

Notes: Bootstrapped standard errors clustered by county. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Controls include year and county fixed effects, the total number of TAA affected workers and petition denial rate of petitions investigated in a given time period.

Table B. 2: List of routine task indexes

Reference	Index	Description
A. Computerisation		
Autor, D., & Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market.	AD RTI routine	Quantifies the routine task content of an occupation. Routine tasks are defined as those that follow explicit rules and can be carried out by machines. It focuses on indicators such as “repeating the same tasks,” “finger dexterity,” or “structured work” to capture cognitive and manual routine components.
	AD RTI manual	It focuses on tasks related to physical skill or coordination. These tasks are not considered routine in nature.
	AD RTI abstract	It focuses on tasks related to reasoning, problem-solving, creativity, management, or interpersonal interaction. These tasks are not considered routine in nature.
Frey, C. B., & Osborne, M. A. (2017). The future of employment: How sustained	Frey RTI	Machine learning and robotics experts evaluate how automatable certain jobs are using sample of 70 detailed O*NET occupations. Experts judge whether job tasks could, with enough data, be fully automated. Jobs were labeled 1 (automatable) or 0 (not or only partly automatable). These data serves as training data. Then using nine key O*NET variables representing technological bottlenecks grouped into perception/manipulation, creative intelligence, and social intelligence (e.g., dexterity, originality, persuasion, caregiving)-a machine learning model is trained to estimate how automatable a job is. The model then predicted automation probabilities for 702 O*NET occupations.
Mihaylov, N., & Tijdens, K. (2019). Measuring the Routine and Non- Routine Task Content of 427 Four-Digit ISCO-08 Occupations.	MT Rck	Measures of the task content of occupations that are based on the International Standard Classification of Occupations 2008 (ISCO-08). Focuses on routine cognitive tasks such as controlling balance sheets, preparing bills and receiving payments, operating cash registers, operating systems and networks, operating laboratory and office computer equipment, testing, inspection and quality control, making hotel reservations, reading work orders, recording and processing information, documents and records, maintaining databases, preparing tax returns, etc.
	MT RMk	Focuses on routine manual tasks such as a setting up, monitoring and operating stationary machinery and equipment, controlling process start-up and shut-down, making standardized products, assembling prefabricated parts and components, sorting and storing, processing, etc.
	MT NRAk	Focuses on non-routine analytic tasks such as researching, analyzing, forecasting, designing, managing, evaluating, treating patients, interpreting data, using advanced software, drafting agreements, creating art, and applying knowledge

Reference	Index	Description
A. Computerisation		
	MT NRIk	Focuses on non-routine interactive tasks such as advising, consulting, teaching, training, interpretation, representation, etc.
	MT NRMk	Focuses on non-routine manual tasks such as involving craft and handwork, operating non-stationary machines and mobile equipment, driving, cooking, serving food, leaning, hair styling, patrolling, repairing, etc.
	MT RTIk	This is the sum of all previous MT scores.
B. Generative AI		
Eloundou, T., Manning, S., Mishkin, K., & Rock, D. (2023). GPTs are GPTs An Early Look at the Labor Market Impact Potential of Large Language Models	GPT4 RTI	This index represents the beta score based on GPT-4 's evaluations of how suitable tasks are for large language models (LLMs). It measures the potential exposure of occupations to LLMs.
Eisfeldt, A. L., Schubert, G., Taska, B., & Ben Zhang, M. (2024). The Labor Impact of Generative AI on Firm Values	GPT4 human	Similar to GPT4 RTI, but based on human assessments of how suitable tasks are for LLMS.
	GenAI core	It measures whether AI affects core occupational tasks those most essential to a job (according to O*NET). If AI can efficiently perform core tasks, the occupation is more likely to be replaced.
	GenAI Supplemental	It measures the tasks that AI can enhance rather than replace supplemental tasks (non-essential or secondary).
	GenAI RTI	It measures Generative AI Exposure. It sums up exposure scores for all tasks in an occupation. Task scores are: 1 (much more efficient with AI), 0.5 (more efficient only with added software), or 0 (not improvable by AI). It reflects general exposure without distinguishing task types.
Felten, E. W., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal	Share supplemental	Measures the share of AI exposure coming from supplemental tasks. A higher value implies more potential for human-AI collaboration rather than replaceme
	AIOE score	The index is built by connecting 10 specific AI applications with 52 occupational skills from the U.S. O*NET database. Using crowdsourced evaluations the strength of each AI application-skill link is rated. These scores are aggregated at the skill level, then combined at the occupation level by weighting each skill 's exposure by its importance and prevalence within a given occupation.

Reference	Index	Description
B. Generative AI		
Webb, Michael. 2020. The impact of artificial intelligence on the labor market. Unpublished manuscript, Stanford.	Webb RTI	This measure aims to quantify which tasks can be performed by AI. It is built by identifying overlaps between AI-related capabilities described in patents and job task descriptions from the U.S. Department of Labor 's O*NET database. Occupations with a higher share of such overlapping tasks are classified as more exposed to AI.
Pizzinelli, C., Panton, A., Tavares, M. M., Cazzaniga, M., & Li, L. (2023). Labor Market Exposure to AI: Cross-country Differences and Distributional Implications	C-AI	It extends AIOE by introducing an occupational-level index called theta, which captures the potential for AI complementarity. The classification of tasks is based on how likely society is to delegate tasks to AI or the need for supportive technologies in specific settings for example, AI is less likely to replace humans in high-risk environments like aviation or medical diagnosis-and reflect reflect the level of education and training required for a job.
Brynjolfsson, E., Mitchell, T., & Rock, D. (2018, May). What can machines learn and what does it mean for occupations and the economy?. In AEA papers and proceedings (Vol. 108, pp. 43-47). 2014 Broadway, Suite 305, Nashville, TN 37203: American Economic Association.	ML RTI	Quantifies how suitable tasks are for being performed by machine learning (ML) or AI. It includes 23 criteria rated on a 5-point scale. The rubric focuses on the technical feasibility of ML performing a task based on current short-term capabilities.
Nedelkoska, L., y Quintini, G. (2018). Automation, skills use and training. OECD Social, Employment and Migration Working Papers No. 202. OECD Publishing, Paris	PIACC AI	The study uses individual-level data from the first two rounds of the PIAAC survey (2011/2012 and 2014/2015), which is the first international survey to comprehensively assess the tasks employees perform and the skills required. Using PIAAC's task-level data, a latent variable model was estimated at the individual level to calculate the likelihood of computerization/automation, following Frey and Osbone (2017) approach.

Reference	Index	Description
C. Offshoring		
Autor, D., & Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market.	AD offshoring	Occupational tasks that are most suitable for offshoring are those that can be performed without the need for physical proximity to customers or specific workplaces. The measure is based on task descriptions especially those involving physical presence and personal interaction (e.g., caregiving, teaching, or external communication)—using O*NET data and aggregate to OCC1990. Offshorability is calculated focusing on non-face-to-face no on-site tasks
Firpo, S., Fortin, N. M., & Lemieux, T. (2011). Occupational tasks and changes in the wage structure (No. 5542). IZA Discussion Papers.	FLM no face-to-face interactions	It focuses on offshoring as occupations with no face-to-face interaction.
	FLM no on-site work	It focuses on offshoring as occupations with no on-site work.
	FLM no decision making	It focuses on offshoring as occupations with no face-to-face decision making work.
Blinder, A. S., & Krueger, A. B. (2013). Alternative measures of offshorability: a survey approach. Journal of Labor Economics, 31(S1), S97-S128. Chicago	BK offshoring	Similar as the measures above but focuses on both no face-to-face and no-onsite work.
Goos, M., Manning, A., & Salomons, A. (2011). Explaining job polarization: the roles of technology, offshoring and institutions. Offshoring and Institutions (December 1, 2011).	GMS offshoring	Measures the share of offshoring cases per occupation, using the ISCO classification, and data from the European Restructuring Monitor (ERM) of the European Monitoring Centre on Change (EMCC), accounting for time variation.
Dingel, J. I., & Neiman, B. (2020). How many jobs can be done at home?. Journal of public economics, 189, 104235	Tele RTI	An occupation is considered to be teleworkable if all of the relevant tasks can be performed remotely. Occupations that are not teleworkable involve at least one task that cannot be carried out remotely. The task content of occupations is sourced from the work content and work activities modules of the O*NET database
	Tele manual	Each authorhot manually assigned values of 0, 0.5, or 1 to each 5-digit SOC code based on introspection and then averaged judgments to assess the possibility of telework.

C Validating the measure of exposure

For this exercise I use data from Pew Research Center. The 119 wave of the Pew survey was carried out from Dec. 12 to Dec. 18, 2022. A total of 11,004 panelists responded out of 12,448 who were sampled, for a response rate of 88%. The margin of sampling error for the full sample of 11,004 respondents is plus or minus 1.4 percentage points. The sample is a stratified, random sample of households selected from the U.S. Postal Service's Delivery Sequence File. All respondents were offered a post-paid incentive for their participation. Among the panelists who completed the survey, weights are calibrated to align with the population benchmarks.

Questionnaire development and testing was developed and rigorously tested by Pew Research Center in consultation with IPSOS. To ensure high-quality data, the Center's researchers performed data quality checks to identify any respondents showing clear patterns of satisficing. This includes checking for very high rates of leaving questions blank, as well as always selecting the first or last answer presented. As a result, only 8 ATP respondents were removed from the survey dataset prior to weighting and analysis.

From this data I extract questions regarding knowledge and concern about AI technologies, especially those associated with impact on work or self. This stands to reason since I'm interested in labor market outcomes, thus available questions on ethical implications of AI are outside of the scope of this paper. Table C4 presents the questions used. I also create a knowledge index which is equal to the percentage of questions in panel B answered correctly.

First, I analyze the Pearson correlations between the questions of interest (Figure C1). We observe that there is a positive correlation between hearing about AI, using AI and correct knowledge about AI. We also observe that those who are more likely to use Internet and social media, are also exhibit more exposure to and knowledge of AI.

Second, we observe that young individuals and highly educated ones exhibit more exposure to and knowledge of AI (Table C1). Per our previous findings, those who are young are more exposed to ICTs as measured by the AD RTI. Also, those educated individuals that should be more exposed to AI, are less likely to make use of the availability bias heuristic.

Third, we map the measure of Google scholar searches about ChatGPTs to respondents in the Pew survey. Table C2 shows that for individuals with higher ChatGPT scores, the measures of exposure to and knowledge of AI are higher. This indicates that Google search scores can indeed reflect average selection into the treatment, but the response will be moderated by the AD RTI.

Figure C. 1: Correlation matrix between PEW survey questions




Table C. 1: Profiles for PEW survey data

	Outcome				
	Heard about AI	Uses AI	Knowledge about AI	AI hurts me	AI hurts workers
Male	-0.205*** (0.017)	-0.060*** (0.020)	-0.552*** (0.080)	0.031* (0.016)	0.038* (0.022)
A. Age					
18-29	0.044 (0.029)	0.097*** (0.037)	0.861*** (0.134)	0.094** (0.037)	0.100** (0.050)
30-49	0.007 (0.024)	0.129*** (0.027)	0.780*** (0.111)	0.066** (0.028)	0.072* (0.039)
50-64	0.024 (0.026)	0.049* (0.026)	0.530*** (0.116)	0.067** (0.028)	0.075* (0.040)
B. Ethnicity					
White non-Hispanic	-0.060** (0.026)	-0.023 (0.028)	0.212** (0.107)	0.003 (0.025)	0.033 (0.033)
Black non-Hispanic	0.039 (0.035)	-0.010 (0.035)	-0.648*** (0.123)	0.056* (0.031)	0.072* (0.041)
Hispanic	0.028 (0.061)	-0.083 (0.062)	-0.235 (0.206)	-0.045 (0.030)	-0.042 (0.068)
C. Education					
Highs school	-0.055 (0.056)	0.057 (0.062)	0.352 (0.251)	-0.075 (0.079)	-0.019 (0.083)
Some college	0.016 (0.058)	0.094 (0.060)	0.910*** (0.227)	-0.048 (0.079)	0.041 (0.076)
Technical	0.062 (0.063)	0.126** (0.064)	0.934*** (0.236)	-0.104 (0.083)	-0.047 (0.083)
College	0.041 (0.059)	0.148** (0.060)	1.535*** (0.224)	-0.089 (0.078)	-0.075 (0.079)
Graduate	0.124** (0.059)	0.226*** (0.063)	1.700*** (0.225)	-0.088 (0.077)	-0.083 (0.077)

	Outcome				
	Heard about AI	Uses AI	Knowledge about AI	AI hurts me	AI hurts workers
D. Area					
Metropolitan	0.017 (0.026)	0.063** (0.028)	-0.007 (0.107)	-0.031 (0.040)	-0.044 (0.036)
Urban	0.048** (0.022)	0.025 (0.025)	0.204** (0.088)	-0.027 (0.027)	-0.013 (0.026)
E. Industry					
Health care and social assistance	-0.094* (0.053)	-0.036 (0.054)	0.198 (0.202)	-0.098 (0.067)	-0.116* (0.069)
Manufacturing, mining, construction, agriculture	-0.067 (0.055)	-0.053 (0.056)	0.271 (0.217)	-0.083 (0.069)	-0.154** (0.071)
Retail and trade	-0.070 (0.055)	-0.091* (0.054)	-0.009 (0.213)	-0.019 (0.073)	-0.083 (0.078)
Education	-0.043 (0.058)	-0.054 (0.054)	-0.091 (0.203)	-0.097 (0.068)	-0.066 (0.071)
Banking, accounting, finance, real state	-0.000 (0.059)	0.017 (0.061)	0.441** (0.216)	-0.049 (0.072)	-0.024 (0.072)
Transportation	-0.025 (0.073)	-0.047 (0.063)	0.414* (0.231)	-0.036 (0.081)	-0.131 (0.089)
Government	-0.025 (0.060)	-0.084 (0.057)	0.346 (0.220)	-0.148** (0.065)	-0.124* (0.071)
Professional, scientific, technical	-0.001 (0.054)	0.031 (0.059)	0.523*** (0.201)	-0.142** (0.064)	-0.082 (0.073)
Observations	5015	4971	5017	3329	4001

Notes: Standard errors clustered by county. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Controls include year and county fixed effects, the total number of TAA affected workers and petition denial rate of petitions investigated in a given time period.

Table C. 2: Profiles for PEW survey data

	Average	Sd.	Min	Max	N	Quartile of ChatGPT google searches			
						Q1	Q2	Q3	Q4
A. Basic opinion about AI									
Heard about AI	0.26	0.44	0.00	1.00	10996	0.18	0.25	0.26	0.25
Uses AI	0.28	0.45	0.00	1.00	10864	0.25	0.25	0.24	0.30
B. Knowledge									
... of AI in costumer service	0.65	0.48	0.00	1.00	11004	0.60	0.67	0.64	0.73
... of AI in music industry	0.57	0.50	0.00	1.00	11004	0.51	0.58	0.56	0.64
... of AI in email services	0.51	0.50	0.00	1.00	11004	0.45	0.50	0.49	0.56
... of AI in health services	0.68	0.46	0.00	1.00	11004	0.63	0.70	0.68	0.72
... of AI in online shopping	0.64	0.48	0.00	1.00	11004	0.59	0.66	0.63	0.69
... of AI in smart home	0.62	0.48	0.00	1.00	11004	0.58	0.62	0.62	0.68
Knowledge index	-0.18	1.87	-3.36	1.81	11004	-0.44	-0.13	-0.23	0.11
C. Media use									
Has twitter account	0.30	0.46	0.00	1.00	10983	0.27	0.32	0.27	0.23
Uses internet frequently	0.44	0.50	0.00	1.00	10830	0.39	0.48	0.39	0.39

Finally, Table C3 shows that individuals living in areas with high Google search scores for ChatGPT and higher levels of RTI are more likely to have know about AI. We do not observe statistically significant effects for the other outcomes. This provides reassurance that Google searches are associated with knowledge about AI.

Table C. 3: Profiles for PEW survey data

	Outcome				
	Heard about AI	Uses AI	Knowledge about AI	AI hurts me	AI hurts workers
ChatGPT scores at Q1 AD RTI	0.071 (0.103)	-0.180* (0.099)	0.689** (0.322)	0.023 (0.091)	-0.006 (0.119)
ChatGPT scores at Q2 AD RTI	-0.050 (0.144)	-0.121 (0.162)	1.027** (0.432)	-0.068 (0.121)	0.025 (0.149)
ChatGPT scores at Q3 AD RTI	-0.123 (0.181)	-0.025 (0.222)	1.733*** (0.540)	-0.157 (0.169)	0.063 (0.194)
ChatGPT scores at Q4 AD RTI	-0.164 (0.234)	0.002 (0.294)	2.185*** (0.666)	-0.320 (0.223)	-0.053 (0.249)
Observations	1879	1866	1879	1378	1605

Notes: Standard errors clustered by county. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Controls include year and county fixed effects, the total number of TAA affected workers and petition denial rate of petitions investigated in a given time period.

Table C. 4: Questions used from PEW survey data

Question	Options	Variable name	Value
A. General opinion questions			
How much have you heard or read about AI?	A lot	Heard about AI	1
	A little		0
	Nothing at all		0
Overall, would you say the increased use of artificial intelligence (AI) in daily life makes you feel	More excited than concerned	Concerned about AI	0
	More concerned than excited		1
	Equally concerned and excited		0
Just your impression, how often do you interact with artificial intelligence (AI)?	Almost constantly	Uses AI	1
	Several times a day		1
	About once a day		0
	Several times a week		0
	Less often		0
B. Thinking about customer service, which of the following uses artificial intelligence (AI)?			
When playing music, which of the following uses artificial intelligence (AI)?	Answers correctly	Knowledge of AI in music industry	1
	Answers incorrectly		0
When using email, which of the following uses artificial intelligence (AI)?	Answers correctly	Knowledge of AI in email services	1
	Answers incorrectly		0
Thinking about health products, which of the following uses artificial intelligence (AI)?	Answers correctly	Knowledge of AI in health services	1
	Answers incorrectly		0
Thinking about online shopping, which of the following uses artificial intelligence (AI)?	Answers correctly	Knowledge of AI in on-line shopping	1
	Answers incorrectly		0
C. Thinking about devices in the home, which of the following uses artificial intelligence (AI)?	Answers correctly	Knowledge of AI in smart home	1
	Answers incorrectly		0
Over the next 20 years, how much impact do you think the use of artificial intelligence (AI) in the workplace will have on [RANDOMIZE ITEMS]			
You, personally	A major impact	AI impacts me	1
	A minor impact		0
	No impact		0
Workers generally	A major impact	AI impacts workers	1
	A minor impact		0
	No impact		0

Question	Options	Variable name	Value
D. Thinking about the use of artificial intelligence (AI) in the workplace over the next 20 years, what do you think the outcome will be for			
You, personally	AI will help more than it hurts	AI hurts me	0
	AI will equally help and hurt		0
	AI will hurt more than it helps		1
Workers generally	AI will help more than it hurts	AI hurts workers	0
	AI will equally help and hurt		0
	AI will hurt more than it helps		1

D Event study design diagnostics

Table D. 1: Optimal bandwidth selection

	Left of r	Right of r
MSE	15.33	15.33
Bias (b)	26.49	26.49
CER	16.22	15.28
Bias (b)	25.26	25.20
MSE-2	10.79	10.79
Bias (b)	26.49	26.49
CER-2	11.42	10.75
Bias (b)	25.26	25.20
Observations	51910	
Fixed effects	Yes	
Controls	Yes	

Note: Controls include: age, gender, race, is migration status, level of education, marital status and has children at home.


D.1 Robustness

Three relevant concerns may arise from our empirical design: i) The possibility that any group of observations is driving the results. ii) The possibility that confounders that are not accounted may overturn the results. I explore these potential issues next:

Parameter stability to observations. To further corroborate that there's no one observation driving the results, I carry out a robustness tests wherein I drop on congressional district at a time with replacement (*à la* Jackknife). I find that my treatment indicator is quite stable and statistically significant for each permutation (Figure D2).

Sensitivity to unobserved confounding. I check the sensitivity of the estimated results with respect to deviations from the conditional exogeneity assumption; i.e., if there are unobserved variables that affect assignment into treatment and the outcome variable simultaneously that estimated coefficients may not be robust to. I explicitly relax the exogeneity assumption by allowing for a limited amount of correlation between treatment and unobserved components of the outcomes (Imbens, 2003). [To be finished as I don't have yet enough confounders to make a meaningful analysisl.]


Figure D. 1: Manipulation checks about the optimal bandwidth for collective bargaining agreements




Note: 95% confidence intervals are clustered at the DMA level and controlling for family-wise error. Controls include: Covid-19 recovery and nightlights. Regressions use observations in a 85-day optimal bandwidth about the release of ChatGPT.

Figure D. 2: Parameter stability to excluding one district with replacement


(a) Number of collective bargaining cases




(b) Bargaining unit size



(c) Number of days from initiation to expiration



(d) New bargaining process



Source: Author's calculations.

Note: Note: 95% Conley confidence intervals in parentheses. Estimates are doubly robust. Controls include: nightlights as a proxy of economic activity, the number of covid cases, and public sector status. Regressions use observations in a 30-day optimal bandwidth about the release of ChatGPT..

E Descriptives

Table E. 1: Summary statistics

	Average	Sd.	Min	Max	N	Quartile of ChatGPT google searches			
						Q1	Q2	Q3	Q4
Google search index for ChatGPT	13.93	20.97	0.00	93.00	35164.00	0.00	7.44	21.69	34.76
Collective bargaining									
Number of cases	1.38	1.69	1.00	62.00	43421.00	1.40	1.67	1.13	1.24
Bargaining unit size	239.59	3553.02	0.00	340000.00	43338.00	168.50	124.95	287.28	171.68
Number of days	105.95	1962.12	0.00	365307.00	41012.00	91.19	88.35	98.73	100.64
Initiates after ChatGPT	0.08	0.27	0.00	1.00	44977.00	0.04	1.00	1.00	1.00