The Electoral Consequences of AI Adoption: Evidence from US Elections*

Alexis Antoniades† Carlos Felipe Balcazar‡
Emmanouil Chatzikonstantinou§ Andreas Kern¶

December 27, 2025

Abstract

What are the impacts of AI adoption on electoral outcomes? Unlike previous research on the electoral consequences of automation, we show that AI can benefit the political left. This occurs as a result of both *economic voting* and *composition* effects in response to AI adoption. Using data from the U.S. presidential elections (2020 and 2024) and a novel measure of AI labor market tightness and a shift-share design, we show that: i) The vote share for Democrats increases as voters swing to the left, reflecting improving local economic conditions. This economic voting effect is driven by ideological moderates. ii) AI adoption manifests in job creation, which attracts more educated workers, who are more likely to be left-leaning. This composition effect also benefits Democrats. We show that AI as not just an economic force but also a political one, shaping electoral outcomes through local labor markets.

Words: 8181

*The comments of Sandy Handan, Michael Becher, and participants at Georgetown University's Political Economy Seminar and AI-PSR at University of Barcelona, are especially acknowledged.

†School of Foreign Service, Georgetown University Qatar. Contact: aa658@georgetown.edu

‡Department of Political Science and School of Public Policy, University College London. Contact: c.balcazar@ucl.ac.uk

§School of Foreign Service, Georgetown University Qatar. Contact: ec1269@georgetown.edu

¶McCourt School of Public Policy, Georgetown University. Contact: ak679@georgetown.edu.

1 Introduction

Artificial Intelligence (AI) is leading the technological transformation of societies. In the U.S. alone, more than one-third of all firms have integrated AI into their workflows; furthermore, over 300 million people worldwide use AI for various tasks.[1] While policymakers see great economic potential in these new technologies, critical voices emphasize the downside risks associated with broader AI adoption. Prime among these are concerns about job displacement, as some firms have become more reluctant to hiring entry-level positions due to AI (Brynjolfsson, Chandar and Chen, 2025). Nevertheless, there is to date no systematic evidence of widespread job loss (Gimbel et al., 2025). Extant research shows that AI adoption has so far been accompanied by increased firm productivity, economic growth and demand for labor (Aghion and Bunel, 2024; Andreadis, Chatzikonstantinou, Kalotychou, Louca and Makridis, 2025; Acemoglu, 2025). Historically, however, socio-political transformation ensues following widespread automation (Boix, 2019).

In Political Science, scholars show that automation in the form of robot adoption has led to increased support for the far-right and economic nationalism (Gallego and Kurer, 2022); robotization does not generate a "rage against the machine" (Wu, 2022; Chaudoin and Mangini, 2024); this is no different for AI (Zhang, 2023; Menon and Zhang, 2025). Recent work has focused on isolating the informational impact of AI adoption on workers' political preferences (e.g., Magistro et al. 2025; Borwein et al. 2025), suggesting that some workers feel economically insecure due to AI. Nonetheless, it remains unclear whether such informational effects translate into meaningful electoral or policy changes. Notwithstanding the evidence, we still know little about the impacts of AI on political behavior. Observational evidence, while important, is scant.

We also observe that—unlike robot adoption—AI adoption is correlated with higher vote shares for democrats in numerous counties (Figure 1). This stands in contrast with the extant literature, which suggests the opposite relationship with automation. Herein, we tackle this puzzle.

The positive correlation in Figure 1 may be because technological innovations such as AI give rise to two competing effects (Aghion and Howitt, 2023): On the one hand, a positive income effect, which captures the gains from higher productivity and profits associated with AI adoption. These gains, in turn, can stimulate job creation. On the other hand, increased automation introduces a substitution effect, as firms may replace workers with machines or algorithms, leading to potential job losses. At the same time, however, the income effect of automation can also reinstitute labor, thereby offsetting the substitution effect to some extent (Acemoglu and Restrepo, 2020). The existing evidence suggests that, at least in the early stages of AI adoption, the income effect tends

[1] See https://explodingtopics.com/blog/ai-statistics and https://explodingtopics.com/blog/companies-using-ai.

to dominate (Andreadis, Kalotychou, Chatzikonstantinou, Louca and Makridis, 2025).

Figure 1: Counties with more liberals have more AI jobs

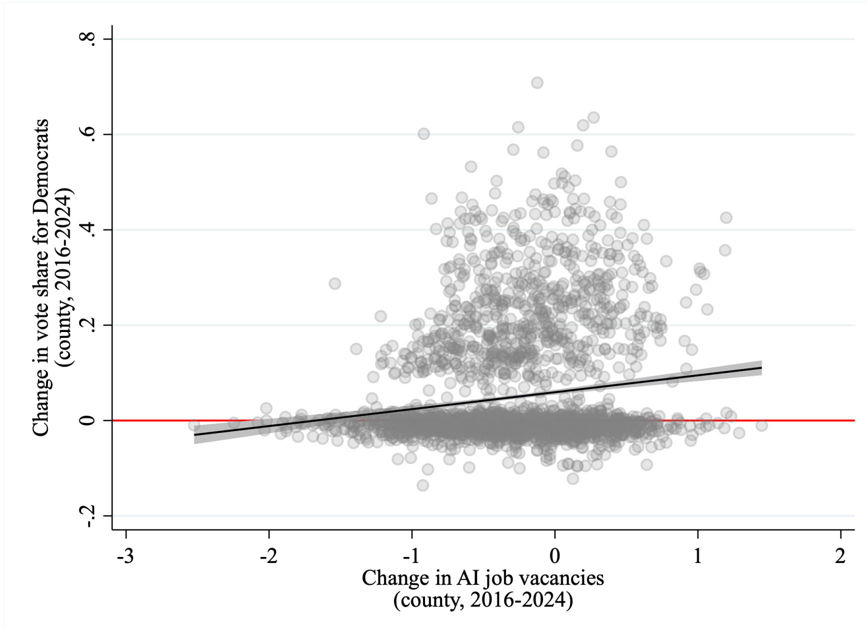

Notes: Data on electoral outcomes obtained from Fox, the New York Times, and David Leipzig's Atlas; data on AI adoption obtained from Lightcast. (See Section 3 for data description.) The figure adds a linear regression line with 95% confidence bands to show the direction of the (positive) correlation. We observe two cluster of observations, one with a positive correlation driving the correlation and one with no correlation at around zero in the x-axis.

In this paper, we examine how local exposure to AI shapes electoral behavior. We combine data on the geographic distribution of AI-intensive occupations and industries with detailed information on voting patterns and political preferences for the US. Specifically, we investigate how changes in labor demand driven by AI adoption translate into voting outcomes. In this way, we provide one of the first systematic analyses of the observed political consequences of AI.

Unlike previous findings showing that automation, particularly robot adoption, tends to shift voters toward the (far-)right, we hypothesize that AI adoption can increase support for the left. The key distinction lies in the nature and distribution of economic gains generated by AI, which differ markedly from earlier waves of automation.

First, AI brings substantial economic benefits. Initial investments in AI infrastructure lead to a local uptick in investment and employment (Fang and Greenstein, 2025). As firms integrate AI into their production processes, demand rises for AI-related occupations and complementary skilled labor. Within adopting firms, AI deployment enhances worker productivity and enables firms to realize significant efficiency gains through process optimization (Babina et al., 2024). Together, these effects translate into tangible local economic gains and sustained increases in demand, especially for highly educated workers (Antoniades, Chatzikonstantinou and Ahman, 2024; Antoniades, Chatzikonstantinou and Savka, 2024), a group that tends to lean left politically.

Second, while these benefits are likely to be concentrated in urban areas, where left-leaning parties already enjoy an electoral advantage (e.g., Mettler and Brown 2022), the political implications of AI are not limited to these strongholds. If AI-driven growth improves local economic conditions through both spillovers and economies of scale (e.g., Andreadis, Kalotychou, Chatzikonstantinou, Louca and Makridis 2025), even moderate or right-leaning voters may reward the incumbent, especially moderates who are more likely to reward good economic performance—a mechanism consistent with extant theories of economic voting (e.g., Kayser and Wlezien 2011; Eggers 2014). In contrast, left-leaning voters and staunch right-wing ones are likely to maintain their partisan alignment regardless of economic conditions. Hence, Democrats can benefit from economic voting even in Republican-leaning areas.

Our argument aligns with standard models of Bayesian updating, in which voters revise their political priors when economic outcomes exceed expectations (Healy and Malhotra, 2013). We refer to this as the *economic voting effect.* However partisanship may moderate this effect, making it more difficult for voters with strong opposing priors to update their beliefs (Bailey, 2019). This implies that the extent of economic voting may depend on the characteristics of the electorate, such as their ideology, and not just government performance (Kayser and Wlezien, 2011). Thus, economic voting effect should be driven by moderates.

In addition, AI may indirectly benefit the left through a *composition effect*: if the growing demand for skilled labor attracts or empowers more highly educated, left-leaning workers in AI-intensive areas, the partisan balance of the electorate itself shifts ceteris paribus.

Taken together, these mechanisms—the economic-voting and composition effects—imply that AI adoption is more likely to benefit the left, widening its electoral advantage in areas where AI-related industries expand most rapidly. The right, in contrast, would gain electorally only if the economic effects of AI are perceived as negative among liberal or moderate voters—consistent with previous work on automation and politics—or if a positive economic voting effect can be attributed to them and it dominates the composition effect.

From here onward we establish a scope condition: There is mixed evidence about the existence of these types of mechanisms in down-ballot races and primaries as opposed to presidential ones (Guntermann, Lenz and Myers, 2021; Holliday, 2022; DeLuca, Moskowitz and Schneer, 2024). This is often attributed to the lower salience of and the higher opportunity cost of voting in down-ballot races, which observe lower participation of moderates and higher participation of ideologues. Therefore, we focus on presidential races because we should observe more meaningful variation in electoral outcomes as a result of economic shocks.

Because our study captures the early stages of AI diffusion, the balance between labor demand and supply plays a central role in amplifying both channels. If labor demand increases as a result of rapid economic expansion due to AI adoption, such conditions strengthen the economic voting mechanism, as voters reward the incumbent president for perceived economic success. At the same time, higher labor demand benefits highly skilled, urban, and left-leaning workers, who take advantage of better economic opportunities, reinforcing the composition effect by increasing their relative presence in AI-intensive areas. Conversely, if labor markets remained slack, weak labor demand would lead to wage stagnation, frustration, and potential backlash against incumbents.

Synthesizing these insights, we interpret excess labor demand from AI adoption primarily as a sign of local economic expansion and favorable perceptions of economic performance. Under these conditions, AI adoption is expected to benefit incumbents and, given the demographic profile of AI-intensive areas, to help the political left who had the incumbent president during the AI boom starting in 2022 with the advent of ChatGPT.

We investigate the hypotheses above by exploiting the quasi-exogenous variation of AI adoption using a shift-share design (Borusyak, Hull and Jaravel, 2025). Specifically, we investigate the impact of AI adoption on the US electoral outcomes between the 2020 and 2024 presidential elections, when the incumbent was a Democrat, using newly collected data on millions of job vacancies from Lightcast—a leading source of labor market information with a vast repository of millions of job postings.[2] We use the *AI shock* driven by the unforeseen advent of commercial and scalable AI technologies, which started with the release of ChatGPT in November of 2022 (Balcazar., 2025), as the shift. The share is defined by the occupation and county. We use the 2016 and 2020 presidential elections as a placebo, as AI had not been widely adopted then.

We measure the concept of excess labor demand using *job market tightness*—or the number of vacancies per unemployed—as is the standard measure for capturing this concept (Barnichon and Shapiro, 2024). Specifically, we focus on the share of job market tightness attributable to AI jobs.

First, we find that an increase in one standard deviation in job market tightness due to AI—about an increase in the share of AI vacancies of 0.1%—increases votes for democrats by 3.4 percentage points (PPs). We also observe an increase in the likelihood that the democrat flips the county by 6.4 percentage points, and increases the margin of victory in the county by approximately 0.5 percentage points. These results indicate that the AI shock helped Democrats in the presidential elections in Republican-leaning areas, and widened the vote share gap in Democratic strongholds.

[2]https://lightcast.io

We also find evidence for both the economic voting and composition effects as follows. First, we find that in Republican-leaning areas the vote share for Democrats increases by about 5 percentage points in response to a 0.1% increase in job market tightness due to AI. Second, we observe a similar positive effect on the vote share for Democrats in Democratic strongholds, amounting to about 2 percentage points. We find that this latter effect is arguably driven by a higher share of highly educated workers, who are more likely to vote for Democrats.[3]

To evaluate these mechanisms, we use data from the Congressional Election Study (CCES). Our results show that increased support for the Democratic incumbent, Joe Biden, is primarily driven by moderate voters. These respondents also report lower unemployment and higher approval of the economy, suggesting that improved local labor market conditions helps explaining this pattern. We find similar results when examining party affiliation instead of ideological self-identification. Using panel data from CCES, we also find support for the composition mechanism.

To mitigate concerns that our results might be driven by competing mechanisms, we implement a series of robustness checks. Importantly, accounting for robot adoption, import competition, Covid 19, and changes in broad local socioeconomics, as well as other unobservable confounders via sensitivity analysis, our results remain intact and unaffected. Following the method used in Balcazar. (2025), we also examine Google searches to rule out that our findings might be driven by potential informational effects. Therefore, our outcomes are unlikely to be explained by the degree of knowledge about AI, but by labor market dynamics. We also show no one county drives our results via Jackknife.

We contribute to multiple strands of academic literature. First, we add to the growing body of research on the political implications of AI adoption. A key innovation of our study lies in its observational design, which distinguishes it from experimental survey-based approaches (e.g., Borwein et al. 2024; Magistro et al. 2025; Menon and Zhang 2025). Our work closely aligns with observational studies, such as that of Balcazar. (2025) and Inglese (2025), which treat the introduction of ChatGPT as a quasi-natural experiment. However, our study differs from this line of work in that we capture AI adoption by firms, allowing us to assess its direct impact on local labor demand and exposure. Using the share of job postings referencing AI as a direct indicator of AI adoption in local labor markets, we introduce a novel measure of local AI exposure. Moreover, by measuring AI exposure through job postings, we examine how changing demand for job skills impacts voting behavior.

Second, we contribute to the literature on economic shocks and their political consequences.

[3] While education increases downstream income, which may cause voters to support lower taxation and thus the right, schooling may increase exposure to liberal values, increasing support for liberals. Evidence shows the latter effect dominates (Scott, 2022).

Previous research has shown that adverse shocks such as import competition, offshoring, and automation tend to benefit right-leaning parties by reducing labor demand and displacing workers (Center, 2022; Gallego and Kurer, 2022; Baccini, Ciobanu and Pelc, 2024; Wu, 2018; Anelli, Colantone and Stanig, 2018; Gallego et al., 2021). Here, our work closely maps on existing studies analyzing the impact of changing economic conditions on presidential elections—see Jensen, Quinn and Weymouth (2017).

In contrast to this work, our study examines the early stages of AI adoption, a period defined by rising productivity, new investment, and increasing demand for skilled labor. We therefore focus on the income effect, in which economic expansion rather than displacement drives political behavior. As such, our work is tangentially related to existing literature studying the political implications of economic booms and resource windfalls owing to them (Guntermann, Lenz and Myers, 2021; Lepers, 2023; Finseraas and Nyhus, 2025).

Translating this logic into the context of AI, we argue that AI operates through changes in labor market tightness, which captures how local labor markets absorb new labor demand. We argue that tighter labor markets strengthen the economic voting mechanism, as voters reward incumbents for improved economic conditions, and they reinforce the composition effect, as highly educated and urban workers, who are more likely to support liberal parties, gain economic advantage. Our findings thus extend research on the political effects of economic shocks to a context in which technological change generates opportunity and growth rather than decline, showing that AI adoption can increase support for liberal candidates, particularly among moderates, when local labor markets benefit from its expansion.

2 AI, jobs, and political behavior

Technological advances in large language models (LLMs) and their derivatives have become a major driver of transformation across multiple industries, fueling automation in both manufacturing and services. The launch of ChatGPT as a consumer-facing AI application in October 2022 marked a significant milestone, making these advancements accessible to a broad audience. ChatGPT alone has over 400 million weekly active users and continues to grow. According to McKinsey & Company, more than 40% of global businesses now use AI technologies; in the United States, over one-third of all companies do so; in fact, tools like ChatGPT are so embedded in society, that "ChatGPT" is consistenly one of the top ten terms searched in Google worldwide.[4]

[4]See https://explodingtopics.com/blog/ai-statistics, https://explodingtopics.com/blog/companies-using-ai, https://backlinko.com/google-searches.

A central question in understanding the political implications of AI is whether its productivity benefits will be broadly shared or concentrated among specific societal groups (Baccini, Ciobanu and Pelc, 2024). In other words, will AI create both winners and losers, thereby shaping political preferences and electoral outcomes?

Existing research approaches this question primarily through the lens of earlier technological and trade shocks, such as outsourcing, ICTs, and robotics, and their influence on political behavior. A substantial body of work demonstrates that these shocks have reduced labor demand, displaced workers, and triggered political backlash that often benefits far-right parties and candidates. This framework centers on what is commonly referred to as the substitution effect in production processes, where technological change replaces human labor, thereby eroding employment opportunities.

In the case of AI, however, large-scale displacement effects have not yet materialized, and whether they will is still up for debate. Much of the emerging literature on the political consequences of AI instead builds on the notion that AI generates anxiety among workers who fear being replaced (e.g., Borwein et al. 2024; Magistro et al. 2024; Menon and Zhang 2025). Yet observational evidence offers little support for this narrative. If machines were truly the origin of anxiety about widespread job losses, workers might be expected to rage against the machines. Instead, empirical studies find no systematic evidence of such backlash (Zhang, 2022; Menon and Zhang, 2025). For instance, Raviv (2024) demonstrates that elite cues following the release of ChatGPT were rarely associated with concerns about technological unemployment; instead, these findings indicate that innovation and economic development lead to minimal partisan divergence between right- and left-wing voters. Consistent with this, workers continue to attribute their economic insecurity to international trade and immigration rather than to automation, which sustains support for economic nationalism and far-right parties (Wu, 2022).[5] Observational evidence suggests that the young's preferences, if anything, may turn left in response to AI (Inglese, 2025).

While this strand of research emphasizes the potential for substitution and displacement, the dominant public and elite narratives about AI remain largely positive. AI is commonly framed as a driver of innovation, productivity, and competitiveness both economically and strategically (Fast and Horvitz, 2017; Neri and Cozman, 2020; Kelley et al., 2021; Bareis and Katzenbach, 2022;

[5]These findings may emerge because much of the existing research relies on experimental designs that prime respondents to imagine losing their jobs to AI. While these experiments can capture attitudes within a controlled setting (Bansak et al., 2023), Chaudoin and Mangini (2022) shows that such framing can heavily shape responses. This suggests that these results may not translate directly to real-world behavior (Barberis, 2013; Naoi, 2020). For instance, Andrew Yang's 2020 U.S. presidential campaign, which centered on technological unemployment, failed to resonate broadly with voters. Accordingly, the perceived risk of being replaced by AI appears to have limited explanatory power for actual attitudes toward AI (Menon and Zhang, 2024; Balcazar., 2025).

Reauscher, 2022; Roe and Perkins, 2023; Ryazanov, Öhman and Björklund, 2025; Wu, 2024).[6]

In line with these perceptions, a rich literature shows that economic booms and windfalls, often triggered by technological innovations or resource discoveries, can yield substantial political dividends. For example, Fedaseyeu, Gilje and Strahan (2015) studies the U.S. fracking boom and finds that local resource windfalls significantly increased Republican vote shares, doubling the probability of a Republican candidate unseating a Democrat. Their results hold across local and federal elections, suggesting that localized economic upswings can reshape political competition. Similarly, Cooper, Kim and Urpelainen (2018), Gazmararian (2025), and Finseraas and Nyhus (2025) corroborate these findings across different settings, emphasizing the partisan impact of such economic windfalls and their capacity to realign voter preferences.

Translating this logic into the context of AI, we propose a framework that emphasizes the income effect of AI adoption. While AI can replace certain tasks, it simultaneously creates demand for new ones and complements workers in others, thereby increasing productivity, wages, and local income. The overall effect depends critically on the ability of local labor markets to adapt to new skill requirements (Acemoglu, Autor, Hazell and Restrepo, 2022). Where adaptation is high, AI adoption generates job creation, higher wages, and local economic expansion. Where adaptation is low, workers should face barriers to transitioning into new roles, resulting in skill mismatches that create structural unemployment, wage stagnation, and slower growth (Pissarides, 2011).

In short, AI adoption can generate either broad-based prosperity or localized dislocation depending on how effectively workers and firms adjust. This duality renders AI's political consequences inherently ambiguous: in areas where AI-driven growth is inclusive, voters are likely to reward incumbents, whereas in areas where adjustment lags, and benefits remain unevenly distributed, discontent and political backlash may emerge.

2.1 Excess labor demand and job market tightness

To understand the impact of AI, we adopt the task-based approach in Acemoglu and Restrepo (2020). This approach is built on the insight that AI will not impact all working tasks equally. AI can be used to automate production processes, potentially leading to the displacement of workers as machines take over their tasks. In some instances, however, it might complement existing tasks or even lead to the creation of new tasks (Milanez, 2023). To facilitate exposition, we develop and present this task-based approach in graphic form in Figure 2.

[6]When AI narratives turn negative, they typically focus not on employment loss but on existential risks, portraying AI as a potential threat to humanity itself. See Jones (2024) for a related discussion.

We start describing Panel A: To simplify the exposition, we use two types of tasks. Type 1 tasks are those where humans have a clear edge, such as critical thinking or emotional intelligence (Frank et al., 2019). Type 2 tasks could be constructing code, evaluating permutations, etc., which are easier to break into steps and thus to automate. The production possibilities frontier (PPF) is given by the line from $(0,2)$ to $(1,0)$; that is, given the current technology, we cannot produce anything outside of the triangle $(0,0)-(0,2)-(1,0)$. Workers can perform all tasks in the square from the origin to $(W1,W2)$. Tasks in the triangle $(W',W2)-(W1,W2)-(W1,W'')$ are outside of the PPF—thus workers can perform these tasks if technology improves. Workers can only perform tasks in the gray-shaded area.

Note that there is demand for the tasks in triangles $(0,W2)-(0,2)-(W',W2)$ and $(W1,0)-(W1,W'')-(1,0)$, but no supply of labor. These areas, which we color with light blue, are job vacancies. Wages should be high for these tasks because the demand for labor is higher than the supply.

This set-up is without loss of generality. The graphic analysis can accommodate a wide variety of cases with different area sizes and shapes for the tasks machines and workers can perform, as well as for the PPF. Likewise, we can use the area defined by the tasks a person can perform within the PPF to determine their wage. Different returns to labor can be accommodated by adding a third dimension reflecting the marginal productivity of each task in the two-dimensional space in the form of contours, as in Panel B, under the standard assumption that the pay-offs for a task equal the marginal productivity of that task.

In Panel C, we consider the introduction of a machine, like AI, which has the potential of performing all tasks in the square from the origin to $(M1,M2)$, which is represented by the diagonal line pattern. However, it needs a human to perform the tasks in the red triangle $(M',W2)-(W1,M2)-(W1,M'')$. That is, machines create demands for new tasks, such as tasks requiring critical thinking (Agrawal Ajay, 2019). While the machines will reduce the pay-offs of all individuals who can perform tasks in the quadrilateral $(0,0)-(0,M2)-(M',M2)-(W1,0)$, note that they generate pay-offs for those of workers who can perform tasks in the red triangle.

In panel D, we plot tasks of three types of individuals: (A)lice, (B)ob, and (C)harles. The set of tasks that these agents can perform is represented in the areas of the rectangles: Alice can perform the tasks in the pink rectangle $(M',M'')-(M',M2)-(W1,M2)-(W1,M'')$; Bob performs tasks in the purple rectangles from the origin to $(W1,W'')$ and $(0,M'')-(0,M2)-(M',M2)-(M',M'')$; Charles can perform tasks in the yellow quadrilateral $(0,M2)-(0,W2)-(W',W2)-(M',M2)$. AI, in this example, cannot perform the tasks that Charles can; it allows Alice to perform new tasks; and it can perform all of Bob's tasks.

As a result, AI may reduce the demand for labor for people like Bob (i.e., a *substitution effect*), but the extent to which this occurs depends on how costly is to automate Bob (i.e., buying machines). In other words, not all employees like Bob will be automated.[7] While AI can take over some tasks performed by Alice, the demand for people like Alice remains because they can perform the tasks in the red triangle using AI.

The effect on Alice's wage depends on how much AI complements her, because while it takes over some tasks, AI and Alice are complementary for new tasks, which are rewarded, given the contour lines in Panel B. Charles, in contrast, is unaffected because he performs tasks that cannot be automated with the current technology.

It is hard to know *ex-ante* if AI will leave workers worse off because workers like Alice could benefit from AI, even if some workers like Bob are substituted. If AI adoption reflects firms' plans to invest and expand, it can also generate spillovers that grow the demand for people like Bob (a reinstitution effect), hence countervailing the substitution effect of AI (Acemoglu and Restrepo, 2020). The problem emerges if no workers—like Alice—can take over the new tasks, which can translate into high structural unemployment and lower economic growth. However, the evidence does not indicate this, as AI has led to increased productivity, economic growth and labor demand despite it has reduced the demand for some entry-level workers (Aghion and Bunel, 2024; Acemoglu, 2025; Andreadis, Chatzikonstantinou, Kalotychou, Louca and Makridis, 2025; Brynjolfsson, Chandar and Chen, 2025).

We measure excess labor demand by computing the job market tightness with respect to AI vacancies (e.g., Abraham, Haltiwanger and Rendell 2020): A more (less) tight labor market occurs when there are more (fewer) job openings than unemployed workers, generating job vacancies. This situation is measured by the ratio of job vacancies to unemployed workers. In our context, job market tightness will be viewed positively by workers, as AI has led to increased labor demand and economic growth; there is little evidence for skill mismatches. Further, the uneven spatial impact of AI on job market tightness can change patterns of support for political parties and incumbents.

Relevantly, we can use the figures above to understand how labor market tightness changes due to AI: Note that in the absence of workers like Alice, there would be more vacancies (blue triangles in Panel A, plus red triangle in Panel C), although the number of unemployed has not changed (purple area in Panel D). This means that job market tightness increases in the absence of workers like Alice, and that these workers are offered higher wages because their supply is lower than the demand, reflecting excess demand for labor.

[7] This can be accommodated through the contours of marginal productivity as in panel b.

Figure 2: Illustration of tasks-based approach

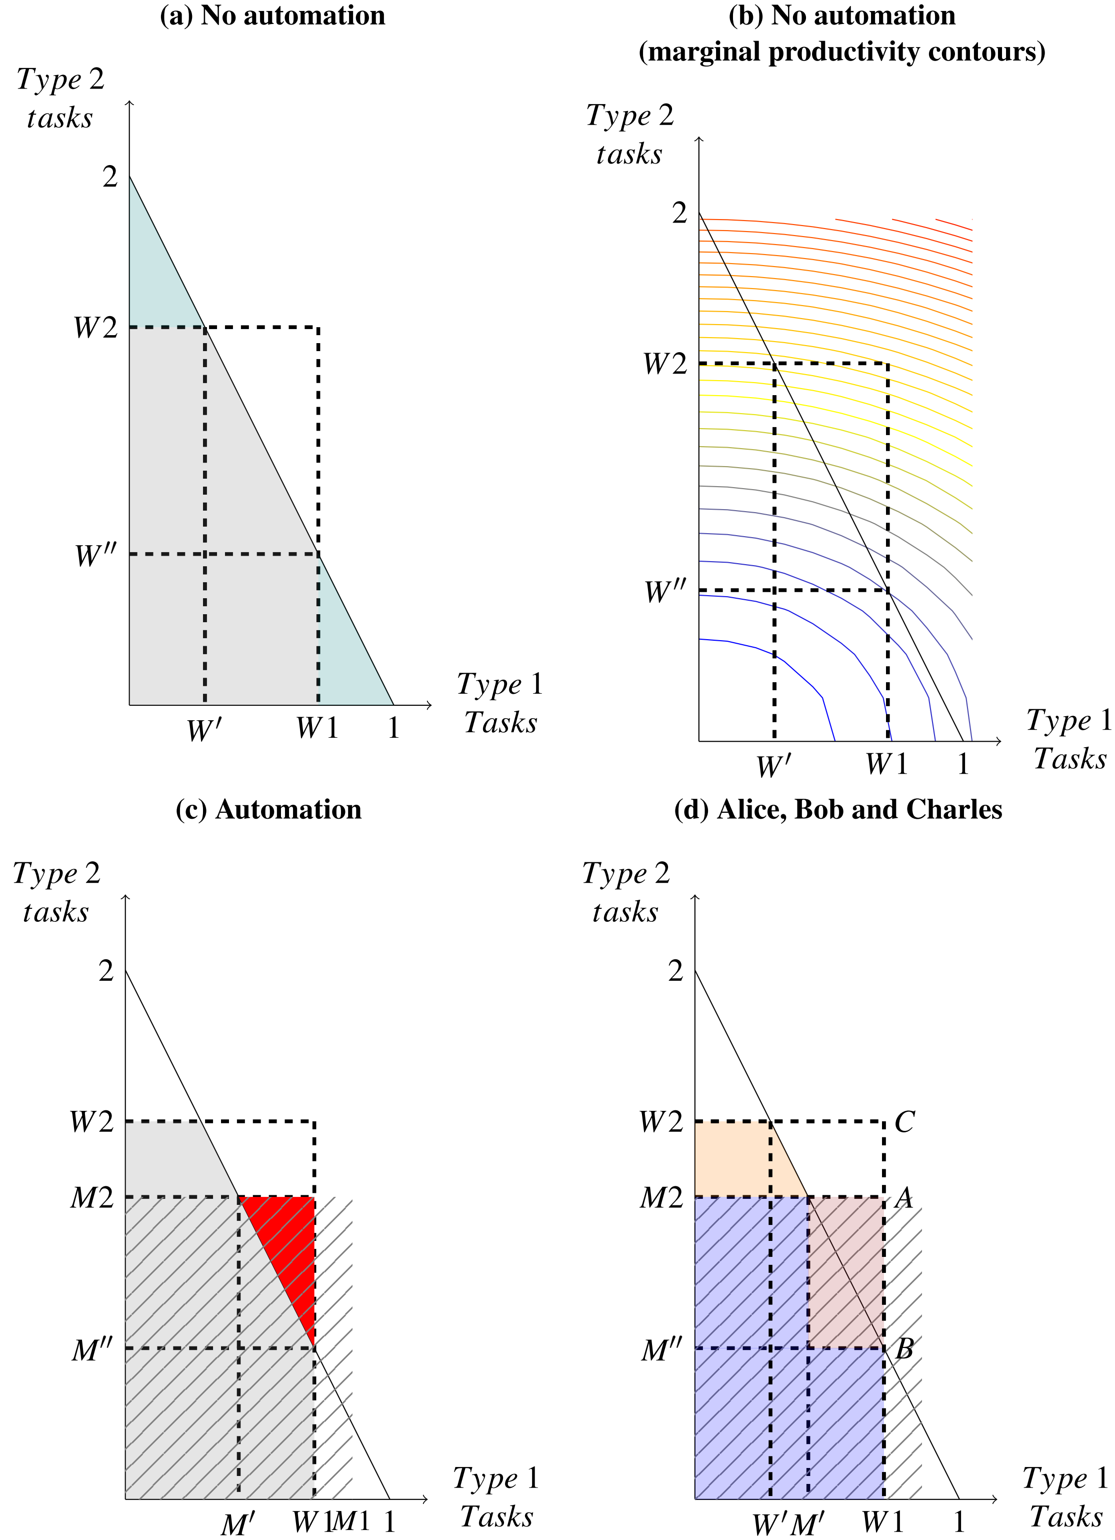

2.2 The impact of AI on politics

We begin by arguing that tightness due to AI is lower in areas with a high supply of workers holding STEM degrees and high levels of education (Antoniades, Chatzikonstantinou and Ahman, 2024; Antoniades, Chatzikonstantinou and Savka, 2024). This occurs because the tasks of AI job vacancies are more likely to necessitate workers with higher education. These areas are also more likely to be urban, where economies of scale are especially conducive to developing and implementing AI (McElheran et al., 2024). Next, building on previous research demonstrating that economic performance is a central determinant of presidential voting behavior (e.g., Lewis-Beck and Tien 2008; Erikson 2009; Jensen, Quinn and Weymouth 2017; Antoniades and Calomiris 2020; Baccini and Weymouth 2021), we argue that these characteristics of AI have relevant implications for politics.

First, educated, urban workers are more likely to vote for liberal candidates because their preferences tend to align more closely with liberal policies, as these individuals are more cosmopolitan (Hainmueller and Hiscox, 2006; Merz, Regel and Lewandowski, 2016; Zingher, 2022). Thus, in areas with low job market tightness, the population may not only be more liberal, but workers, especially the educated, are also more likely to benefit from AI. In these areas, support for liberal politicians is likely to be high to begin with due to the composition of the electorate. For instance, for the US, we observe a positive correlation between the number of AI jobs available and the revealed share of Democrats (Figure 3).

The theory of economic voting posits that voters reward incumbents during good times and punish them during times of economic downturn. Thus, in these areas, we could also observe higher support for the incumbent in response to AI. Nevertheless, we posit some nuances regarding the economic voting effect: If a local voter already votes for liberals, the economic voting effect should solidify her choices. This will not benefit the liberal incumbent because it does not switch a vote in her favor. However, this may benefit the liberal if undecided voters or even conservative ones reward the incumbent for the good economic performance with their vote. This is consistent with the standard Bayesian approach to economic voting theory (Healy and Malhotra, 2013).

This is also consistent with work by Kayser and Wlezien (2011) and Eggers (2014) whereby partisanship reduces accountability and thus economic voting. That is, voters' partisan attachments make them less responsive to government performance, thus voters with weak partisan preference are more likely to do so. For example, Kayser and Wlezien (2011) demonstrates that changing variance in aggregate-level partisan attachment alone is sufficient to alter the strength of the economic vote, because where the variance is low in partisanship, better performance will capture more voters. In other words, economic voting should be driven by moderates because they are

Figure 3: Counties with more AI jobs have more liberals

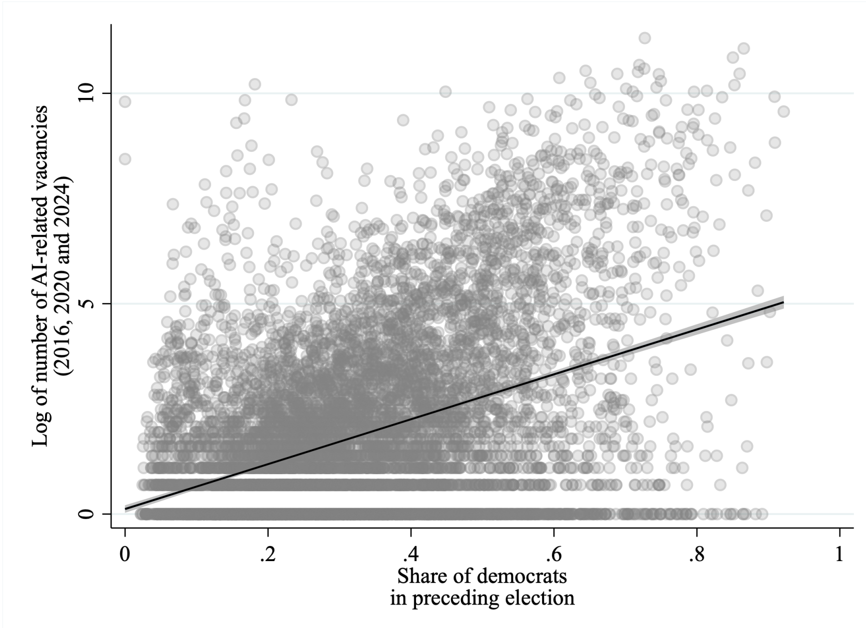

Notes: Data on electoral outcomes obtained from Fox, the New York Times, and David Leipzig's Atlas; data on AI adoption obtained from Lightcast. (See Section 3 for data description.) The figure adds a linear regression line with 95% confidence bands to show the direction of the (positive) correlation.

more likely to update their priors (Bailey, 2019).

Table 1 exemplifies the logic, considering that the degree to which voters update their priors about a politician or party may depend on partisanship. We take into account, however, that liberals are more likely to vote for the left, and conservatives for the right; moderates are likely to be independent (e.g., Gerber, Huber and Washington 2010). Thus, while the relation between partisanship and ideology is not one-to-one, support for left- or right-wing politicians aligns with ideology. Thus, we use ideology and party affiliation interchangeably for the remaining of the section.

Table 1 indicates that moderates should be more likely to update their priors. Liberals may also be more likely to support left-wing politicians due to their good performance, whereas conservatives may not be initially inclined to do so, given their prior beliefs. Likewise, liberals may also be more likely to punish left-wing politicians for poor performance, as they update their priors on the co-partisan. In contrast, conservatives may do so from the outset, given their initial priors. A similar logic applies to right-wing politicians.

Left-wing incumbents are most likely to be rewarded—especially by moderates—in areas where AI adoption has led to high labor market tightness. This effect arises for two reasons. First, AI-driven demand for skilled workers has stimulated employment and income growth, so voters are likely to reward incumbents for good economic performance. Second, in left-leaning ar-

Table 1: Economic voting effect of AI shock on support for Liberal incumbent

Effect of AI	Liberal	Moderate	Conservative
Positive	No effect	Reward	Little to no effect
Negative	Little to no punish	Punish	No effect

eas, where partisan loyalties are already established, such gains mostly reinforce existing support rather than shifting votes. Consequently, left-wing incumbents should perform best in competitive areas, where voters are more responsive to local economic conditions. Right-wing politicians (who are not incumbents in this example), by contrast, benefit primarily when AI adoption leads to job losses or when labor market tightness remains low.[8]

Labor-market tightness in AI also reshapes the political landscape through its effect on population composition. Areas with sustained demand for highly educated and technical workers tend to attract liberal, urban-oriented populations, further aligning AI-intensive areas with liberal electorates—see Berriochoa (2025) and Dasgupta and Ramirez (2025). This composition effect benefits the left more than right, as migration and occupational clustering strengthen the link between AI-driven labor demand and liberal voting behavior.

Overall, variation in labor-market tightness in AI influences politics through two complementary mechanisms: an *economic voting* effect reflects how voters respond to prosperity associated with AI-driven demand, while a *composition effect* captures how the geography of these opportunities reshapes the partisan and ideological map. Together, these dynamics help explain why areas with greater labor-market tightness tend to show stronger support for left-wing candidates—especially if they are incumbents.

3 Data

We focus on the United States as our case study. On the one hand, AI adoption is high in the US, with about one-third of US companies having deployed AI into their activities.[9] Recent research also indicates that AI has had an impact on the labor market (e.g., Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo and Zolas 2022; Aghion et al. 2022; Hampole et al. 2025; Demirci, Hannane and Zhu 2025). On the

[8] The opposite may occur in another context where labor market tightness reflects an inefficient allocation of resources through both skill mismatch and structural unemployment.

[9] See `https://explodingtopics.com/blog/companies-using-ai`.

other hand, the wealth of sub-national data and data on public opinion allows us to evaluate the potential mechanisms connecting AI with voting outcomes more accurately.

3.1 Main dependent variable: Electoral Outcomes

We investigate the impact of AI deployment shifts on voting for Democrats—the left-wing party—in the presidential election. We focus on presidential races because they are higher salience and have higher participation by moderates, in contrast to down-ballot races where ideology itself plays a bigger role (Guntermann, Lenz and Myers, 2021; Holliday, 2022; DeLuca, Moskowitz and Schneer, 2024). Hence, we have more meaningful identifying variation for our purposes using presidential-election outcomes.

Given that AI deployment is a relatively recent phenomenon, we focus our analysis on the 2016, 2020, and 2024 presidential elections. We elaborate on this choice further below. Vote counts for 2016 and 2020 are obtained from MiT election lab; for 2024 the data is obtained from Fox, the New York Times, and David Leipzig's Atlas.[10] With this data we construct measures of the vote share for democrats in the presidentials at the county level.

3.2 Main Independent Variable: Job market tightness due to AI

We construct a shift-share measure that leverages local differences in exposure to national AI adoption trends. Shift-share designs have been widely used to study the effect of automation on a wide variety of outcomes (Borusyak, Hull and Jaravel, 2025).

Specifically, we define

$$\Delta AI_{ct} = \sum_j \left(\frac{AI_{c,j,t}}{V_{c,j,t}} - \frac{AI_{c,j,t-1}}{V_{c,j,t-1}} \right) \times \frac{E_{c,j,t_0}}{E_{c,t_0}}, \tag{1}$$

where c denotes the county, j is the occupation, and t the time period. This measure combines county-level variation in the number of AI vacancies ($AI_{c,j,t}$) over the total number of vacancies ($V_{c,j,t}$)—the *shift*—and baseline employment shares $E_{c,j,t_0}/E_{c,t_0}$ in county c—the *share*.

[10]Fox: https://feeds-elections.foxnews.com/archive/politics/elections/2024/3/2024_Generals/President/[state_abbr]/county_level_results/file.json; New York Times: https://static01.nyt.com/elections-assets/pages/data/2024-11-05/results-[formatted_state]-president.json; David Leipzig's Atlas: https://uselectionatlas.org

Notice we can rewrite the expression above as

$$\Delta AI_{ct} = \sum_j \left(\frac{AI_{c,j,t}}{S_{c,j,t}} \cdot \frac{S_{c,j,t}}{V_{c,j,t}} - \frac{AI_{c,j,t-1}}{S_{c,j,t-1}} \cdot \frac{S_{c,j,t-1}}{V_{c,j,t-1}} \right) \times \frac{E_{c,j,t_0}}{E_{c,t_0}} = \sum_j \left(\frac{T^{AI}_{c,j,t}}{T_{c,j,t}} - \frac{T^{AI}_{c,j,t-1}}{T_{c,j,t-1}} \right) \times \frac{E_{c,j,t_0}}{E_{c,t_0}} \quad (2)$$

where $S_{c,j,t}$ denotes the number job seekers. Thus $T_{c,j,t}$, for instance, denotes the labor market tightness in period t, and $T^{AI}_{c,j,t}/T_{c,j,t} = 1 - T^{\neg AI}_{c,j,t}/T_{c,j,t}$, which is the share of labor market tightness due to AI vacancies.

ΔAI_{ct} captures changes in the efficiency with which the labor market matches job openings to job seekers, attributable to AI jobs (e.g., Barnichon and Shapiro 2024; Bernanke and Blanchard 2025). Thus, our measure is commensurate with the graphical model above by considering excess AI-related labor demand (Figure 2).

To measure the shares, we use the ACS 5-year sample 2005-2009 and the place of work and (6-digit) occupations variables to count the number of full-time workers in each occupation and each county for the year 2009. Therefore $t_0 \approx 2009$.

We measure AI adoption using mentions of AI-related skills from job postings derived from Lightcast. Lightcast is a leading source with a vast repository of millions of job postings. It builds on extracted information from more than 40,000 online job boards, newspapers, and employer websites. It employs advanced machine learning techniques to streamline the data, eliminating duplicate job postings, whether posted multiple times on the same site or across multiple sites, and enriching profiles using standardized information on job titles, company names, skills, and educational requirements.[11]

This data has numerous advantages: The data contains a comprehensive occupational taxonomy built hierarchically, with over 1,900+ specialized occupations that are mapped to the Standardized Occupation Classification (SOC) 6-digit used in the official publications by the BLS. Lightcast also provides precise location data for job postings, enabling us to capture labor demand at the county level. We use this data to obtain data on the number of vacancies: $V_{c,j,t}$.

Using this data, we also identify the number of AI-related vacancies, that is $AI_{c,j,t}$. The full list of keywords is available in Appendix E. We analyze AI-related job postings by extracting relevant skills and keywords from vacancy descriptions, using the Lightcast Open Skills Taxonomy, which includes over 15,000 skills from 2014-2021.[12] We rely on job postings to count job posts

[11] For more details, see their methodology: https://kb.lightcast.io/en/articles/6957446-job-posting-analytics-jpa-methodology.

[12] Although Lightcast data extend back to 2007, we use data from 2014 onward due to quality concerns in earlier years.

containing explicit AI-related skills for each year, county, and occupation. Then, we quantify AI adoption at the local level by counting the number of vacancies containing explicit AI-related skills for each county in the years before the election. Since AI skills demand may be related with general IT occupations demand unrelated to AI, we narrow the AI measure by excluding computer occupations for robustness checks and counting the adoption of AI in non-IT occupations.

Importantly, in identifying the causal effects of AI adoption, we need either the shift or the share to be exogenous. We adopt the latter approach—see Borusyak, Hull and Jaravel (2022). We do this by exploiting the release of ChatGPT in 2022.

While automation may have been present in the form of robots and ICTs prior to 2022, AI was still an infant for commercial operations at scale: There was relatively low interest in AI until 1993. During the 2000s, AI developments were mostly experimental, like developing Deep Blue, Go, Roomba, pattern recognition, and speech recognition. However, tools like deep neural networks couldn't be used at scale and efficiently due to limitations in computer power, which began to be overcome in the 2010s. Convoluted Neural Networks and Artificial Neural Networks, made possible by powerful GPUs. Further, we exploit the timing of the transformer architecture, which debuted in 2017, which allowed us to produce impressive generative AI applications using large language models (LLMs) that led to the development of scalable AI, such as ChatGPT, Gemini, Claude, and others.[13] The technology was not yet largely commercial nor scalable until the release of technologies like ChatGPT. Balcazar. (2025) shows that the advent of ChatGPT acted like a quasi-unforeseen event—see also Appendix A. This makes it an ideal quasi-experiment.

$\Delta AI_{j,t_0,t_1}$ measures the change in AI vacancies responding to the ChatGPT shock. Therefore, it captures changes in labor demand due to this shock. Figure 4 shows the spatial distribution of the shock. Unlike import competition, off-shoring, and robot adoption shocks, we do not observe a high concentration of the shock in the Rust Belt. In Figure A4 in the Appendix, we show that secular trends owing to import competition, off-shoring, and robot adoption are negatively correlated, and at best uncorrelated, with the AI shock. Antoniades, Chatzikonstantinou and Ahman (2024) and Antoniades, Chatzikonstantinou and Savka (2024) also provide evidence that there is no substantive correlation between AI adoption and other trends, and that the shift-share measure is robust to excluding occupations with the largest Rottemberg weights (Goldsmith-Pinkham, Sorkin and Swift, 2020).

These empirical features of our design imply that we focus on the short-run responses to AI

[13]The transformer architecture converts text into numerical representations using embeddings, via a parallel multi-head attention mechanism. This reduces the training time for large language models, allowing them to handle massive amounts of data efficiently for the development of pre-trained systems, such as generative pre-trained transformers (GPTs).

adoption, holding fixed its long-run adjustment (Jaeger, Ruist and Stuhler, 2018). That is, local shocks may trigger adjustments that gradually offset their local impact to an extent, towards a new labor market equilibrium. In this way, the shift-share captures the short-run effects of AI adoption. Further, it also captures the spatial interdependence through the concept of treatment intake—or dose-response—by accounting for the geographical concentration in AI adoption. The identification assumption rests on the fact that areas adopting more AI do so thanks to technological innovations occurring across occupations, which trigger changes in the choices to automate tasks by firms to remain competitive. That is, our shift-share is a measure of the degree of compliance with the treatment from the compliers.

Figure 4: Spatial distribution of changes in tightness due to AI (2020-2024)

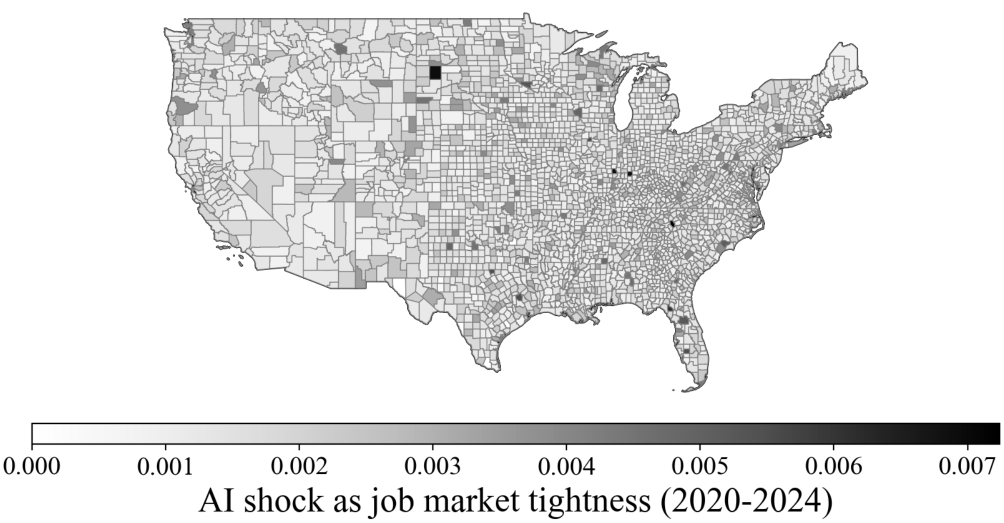

Notes: Job market tightness due to AI is computed as defined in Equation 2.

3.3 Controls

We also obtain data on a number of variables that could affect both AI adoption and political outcomes, like sociodemographics, local economic characteristics, and local labor union strength. Data on population sizes, employment, and other socio-demographics are drawn from the American Community Survey; data on other local industry characteristics come from the Quarterly Census of Employment and Wages and County Business Patterns. Data on pre-treatment exposure to Chinese imports, robots, and deindustrialization come from Baccini and Weymouth (2021) and Balcazar (2023). Data on local union strength is gleaned from the LM form of the Office of Labor-Management Standards (OLMS) (e.g., Becher and Stegmuller 2023). We also control for levels of pandemic recovery rates using data on cases and deaths from the Centers for Disease Control and Prevention in the US and the New York Times.

Table C1 shows the descriptive stats.

4 Research Design

We estimate the following regression equation:

$$\Delta(Y_c;t'_0,t'_1) = \beta \Delta AI_{c,\{t_0,t_1\}} + \delta \Delta D_{c,\{t_0-1,t_1-1\}} + \phi_d + \Delta\varepsilon_{c,t,t-1}$$

where $\Delta(Y_c;t'_0,t'_1)$ denotes the post-treatment change in the outcome of interest at the county level; $\Delta AI_{c,\{t_0,t_1\}}$ is the AI shock as defined earlier; $\Delta D_{c,\{t_0-1,t_1-1\}}$ denotes a vector of local differences in pre-treatment socio-demographic and economic characteristics, like population size, share of those with bachelors degree and above, share of whites, unemployment, manufacturing and agriculture employment and establishments, and many other. We also include lagged measures of the Chinese imports shock, robot adoption, interest in AI via Google searches, and Covid 19. Note that by being in first differences, we already control for county fixed effects; ϕ_d thus denotes congressional district $\times$ year fixed effects, capturing congressional-district-specific time trends; these trends capture the time-variant covariates at the congressional district level.

Our standard errors are clustered at the congressional district level to allow for some degree of spatial correlation, following the recommendations by (Abadie et al., 2023). We use generalized doubly-robust estimators to address the problem of negative weights (Arkhangelsky et al., 2024).

β is our coefficient of interest. It measures the impact of the change in AI vacancies responding to the AI shock. We normalize $\Delta AI_{c,\{t_0,t_1\}}$ to one standard deviation, or about an increase in the share of AI vacancies of 0.1%. β can also be interpreted as the effect of an increase in labor market tightness due to AI—or our measure of excess demand for labor. We anticipate that $\beta > 0$ insofar as AI has led to increased demand for labor and economic growth.

5 Results

We first analyze the impact of the change in AI vacancies responding to the AI shock on the presidential elections. Table 2 shows the results. In Column (1) we observe that an increase in one standard deviation in job market tightness owing to AI increases the vote share for democrats by

3.4 percentage points (PPs); it increases the chances that the democrat flips the county—obtaining a majority of votes—by 6.4 PP (Column 2); it reduces the likelihood of close races by 5.7 PP (Column 3); it increases the margin of victory in the county by about 0.5 PP (Column 4). These results indicate the AI shock helped democrats in the presidential elections.

Although even small shifts in vote share can be decisive in US presidential elections,[14] these effects evidently didn't secure the Democrats' victory; they show instead that AI has impacts on political outcomes. Nevertheless, it is important to put the size of these effects in context. A 3.4 PPs increase in the vote share, for instance, corresponds to one-fifth of a standard deviation in this outcome—similarly for the Democrat flipping the county. For the margin of victory in the county, the effect is about half a standard deviation. Therefore, the effect of the AI shock on voting outcomes is less than proportional.

Table 2: Effect of AI intensity on presidential outcomes (2020-2024)

	Outcomes for democrats			
	Vote share	***Flips county***	***Close race***	***Margin of Victory***
	(1)	(2)	(3)	(4)
AI shock	0.0339***	0.0636***	-0.0566***	0.0054** [t]
	(0.0062)	(0.0204)	(0.0210)	(0.0027)
Observations	2715	2715	2715	2715
Congressional District FE	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes
Industry controls	Yes	Yes	Yes	Yes
Other shocks	Yes	Yes	Yes	Yes
Union power	Yes	Yes	Yes	Yes

Notes: Standard errors clustered by district in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

Next, we investigate if the pre-treatment democrat vote share moderates the effect above. First, Figure 5, Panel A, shows that the vote share for Democrats increases in previously Republican strongholds, where republicans obtained more than 70% of the vote, increases in response to the AI shock by about 5 PP. This implies that an increase in job market tightness is affecting Republicans

[14] See https://www.washingtonpost.com/graphics/politics/2016

through the economic effect because voters are likely rewarding the democrat incumbent at the time—Joe Biden.

Second, Figure 5, panel b, provides nuances to the previous result. On the one hand, it shows that counties flip from Democrat to Republican, especially in the middle of the distribution of pre-treatment Democrat vote shares, in response to the AI shock. In Democratic strongholds, the effects are smaller and non-significant, suggesting that the economic voting effect doesn't dominate in these areas.

Third, Figure 5, panel c, shows that the likelihood of close races falls in Democratic strongholds. Thus, AI seems to be opening the vote share gap in these areas in favor of Democrats.

Lastly, Figure 5, panel d, shows the effect of the AI shock on the margin of victory by democrats. The point estimates are smaller in Republican strongholds.

We also corroborate that these results are not driven by the distribution of the rural population, who leans Republican and thus can be associated to the revealed share of democrats (Figure B5). Therefore, the revealed share of democrats is a distinct moderator from the former.

All in all, these results also point to the presence of a composition effect that is offsetting the weaker economic voting effect in Democratic strongholds. To probe the viability of this mechanism, we estimate the effects of changes in the composition of the population, finding that the proportion of educated individuals increased in areas where democrats obtained a larger vote share in the past but not in Republican strongholds (Figure 6). We do not observe evidence for other mechanisms that may point to other types of composition effects, such as change in the share of white population, a younger workforce, or changes in the prices of housing (Figure B4).[15]

All in all, we observe evidence for an economic voting effect against Republicans in republican strongholds, and evidence of a composition effect in favor of democrats in democrat strongholds. Thus, both economic voting and composition effects benefit democrats.

In Appendix B, we find that our results above are robust to the presence of unobservable confounders via sensitivity analysis. We also show that they are not driven specifically by any given county where treatment intake could be high (e.g., Washington, D.C., or San Francisco). We also measure AI intensity, including software-related jobs, and find that our results hold. Lastly, we use the period 2016-2020 as a placebo test, as AI had not been widely adopted at that time and had therefore not generated a visible effect; we find evidence in support of the placebo.

[15] We further prove the composition mechanism using recently released panel data from the CCES, which tracks if individuals moved congressional districts between survey rounds. We show our analysis in Appendix B.

Figure 5: Heterogeneous Effects of AI intensity on approval for president by revealed share of democrats

(a) Democrat vote share

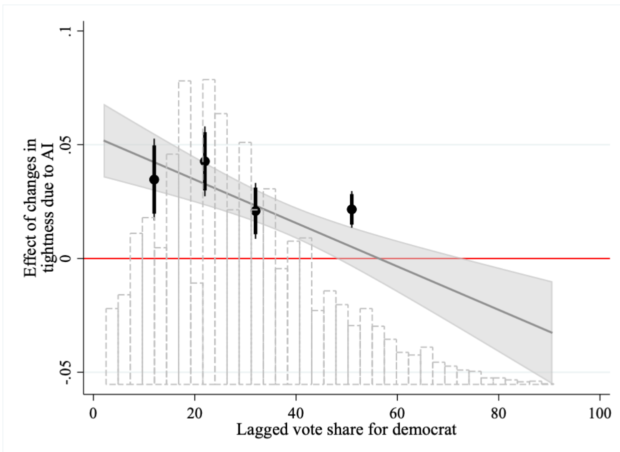

(b) Flips county

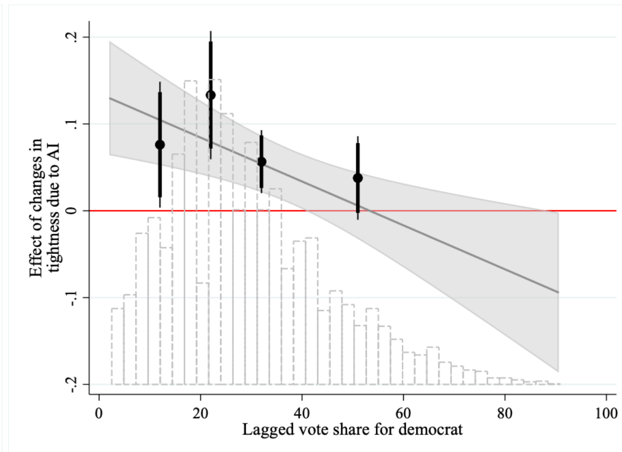

(c) Close race (<=0.05 margin)

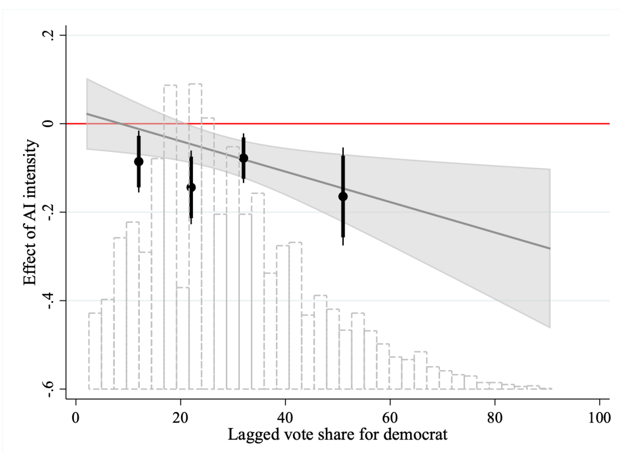

(d) Democrat margin of victory

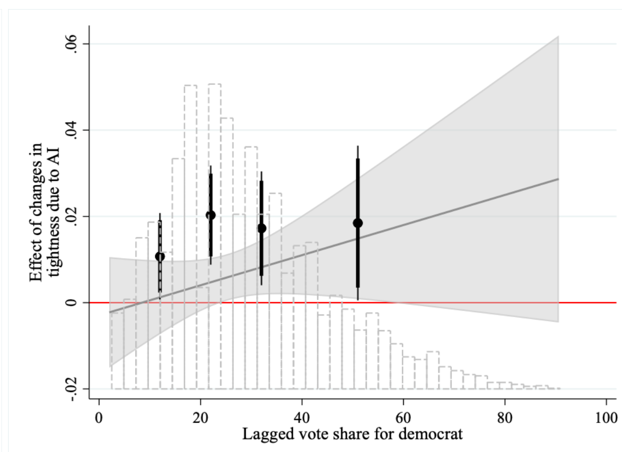

Notes: Confidence intervals and confidence bands are clustered by district. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling for congressional-district fixed effects, which capture location × time non-monotonic trends given that our design is in the first differences.

6 Mechanisms

To explore the key underlying mechanisms at play, we analyze respondents' opinions on the incumbent—Joe Biden. To do this, we use data from the Cooperative Congressional Election Study (CCES), which asks questions about approval for the president and has a large nationally representative sample. We expect our results for public opinion on these selected outcomes to be consistent with our theory. In particular, we analyze the controlled directed effects of AI on party

Figure 6: Heterogeneous Effects of AI intensity on change in share of population with university education

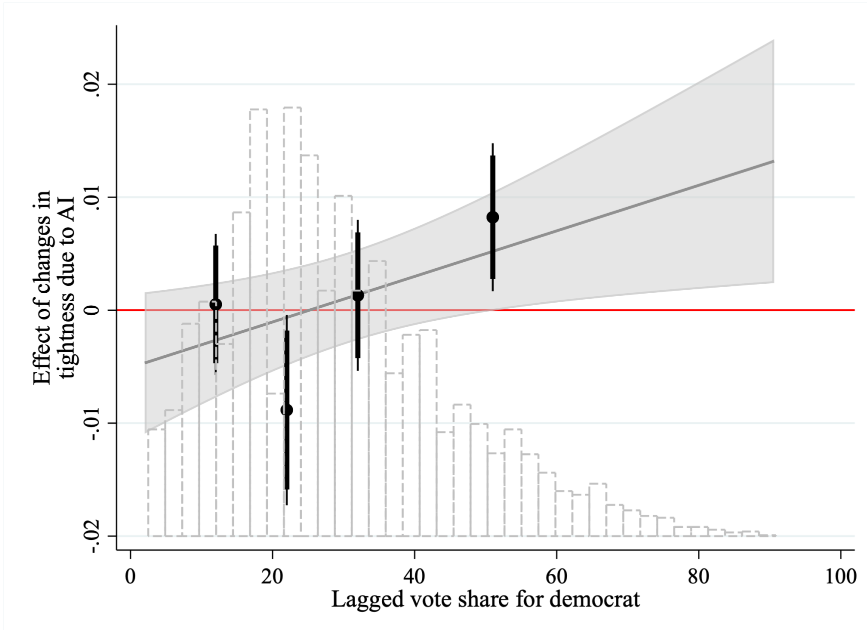

Notes: Confidence intervals and confidence bands are clustered by district. Estimates are heterogeneity robust. Controls include: changes in size of the population, in the share of female labor, Hispanic, white, black and Asian groups, changes in the share of people with high school, college and masters degrees, and in the share of people with 65 years of age and above; changes in the share of manufacturing and light manufacturing in industry; pre-treatment changes in unionization and Right-to-Work Laws; the China import shock, the robot adoption shock, c and other measures of deindustrialization.

identification and to avoid any problems with potential post-treatment bias.

We estimate the following regression equation:

$$\begin{aligned}[\Delta(Y_{i,c};t_0',t_1')|P_i = p] =& [\beta_1 \Delta AI_{c,\{t_0,t_1\}} + \delta \Delta D_{c,\{t_0-1,t_1-1\}} + \delta \Delta X_{i,c,\{t_0-1,t_1-1\}} \\ &+ \phi_d + \Delta \varepsilon_{i,c,\{t_0-1,t_1-1\}} | P_i = p]\end{aligned}$$

where P denotes party and

$$p \in \{liberal,\ neutral,\ conservative\}.$$

or

$$p \in \{democrat,\ independent,\ republican\}.$$

$X_{i,c,\{t_0-1,t_1-1\}}$ includes pre-treatment characteristics of the individual, such as age, gender, and ethnicity, and all other covariates we used in the previous regressions. Our standard errors are clustered at the county level.

Figure 7, Panel B, shows that independents increase their support for the democrat incumbent. Panel A indicates that moderates drive the effect, consistent with our theory. This is plausible insofar as we show in Table C2 that moderates are more likely to lean Democrat than Republican, and to exhibit comparatively higher support for liberal policies v. conservative ones.

Figure 7: Heterogeneous Effects of AI intensity on approval for president by ideology and party identification

(a) By political orientation **(b) By party affiliation**

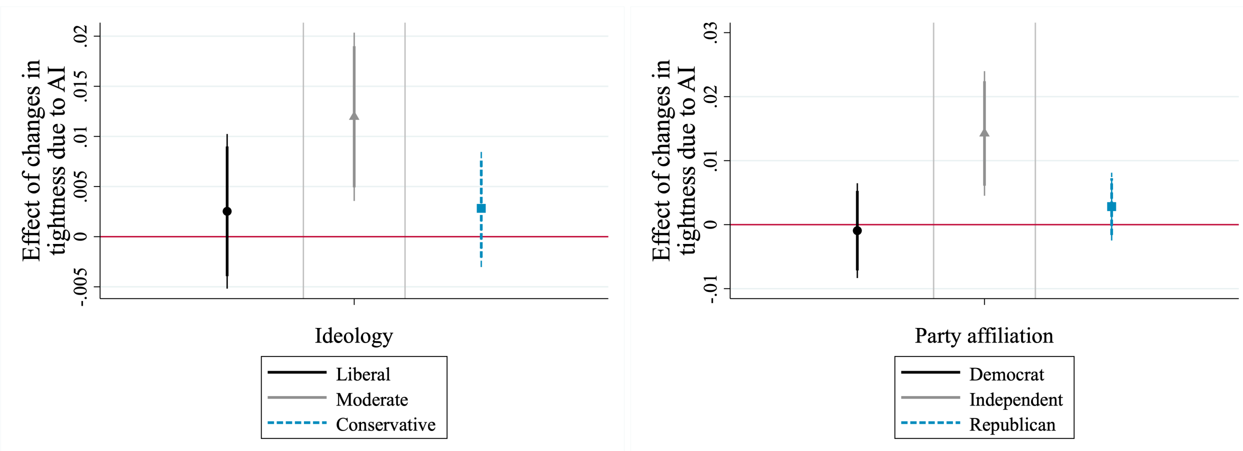

Notes: Confidence intervals are clustered by county. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling for congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

To evaluate the economic voting effect, we use data from the CCES, which asks questions about employment status and perceptions about the economy. We analyze these outcomes following the same design as above.

Figure 8, Panels A and B show the results of our analysis. The results in Panels A and B confirm the notion that democrats are likely to be rewarded by the good labor market, as independents and moderates report a lower likelihood of unemployment; we observe similar results for perceptions about how well the economy is doing. This is consistent with the claim we made above, whereby moderates are driving the effect, and provides additional evidence for the economic voting mechanism.

Figure 8: Heterogeneous Effects of AI intensity on employment status by ideology and party identification

(a) By party ideology, unemployment

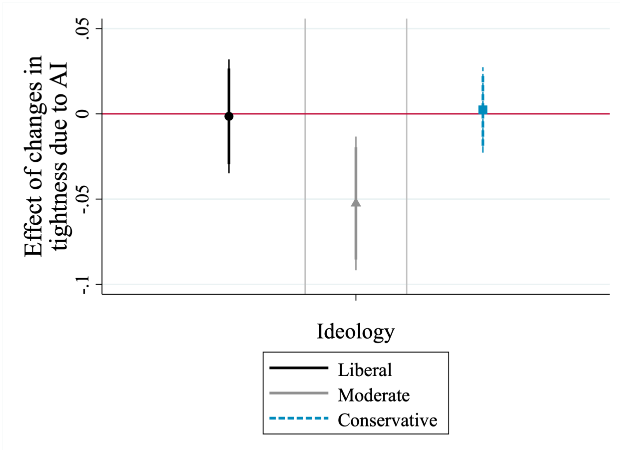

(b) By party affiliation, unemployment

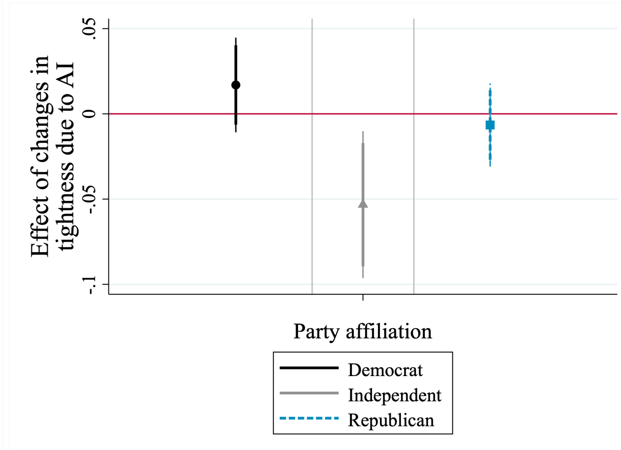

(c) By party ideology, positive opinion about economy

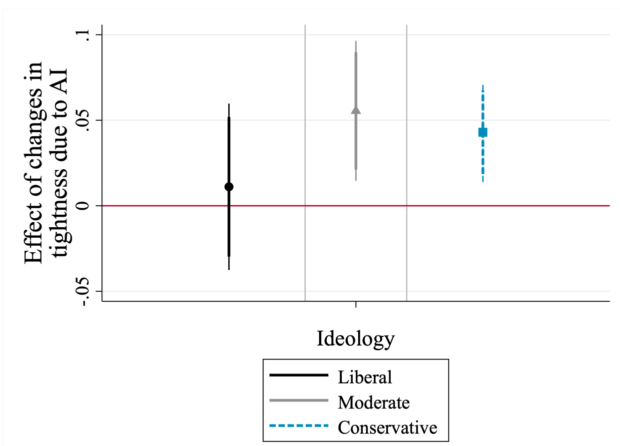

(d) By party affiliation, positive opinion about economy

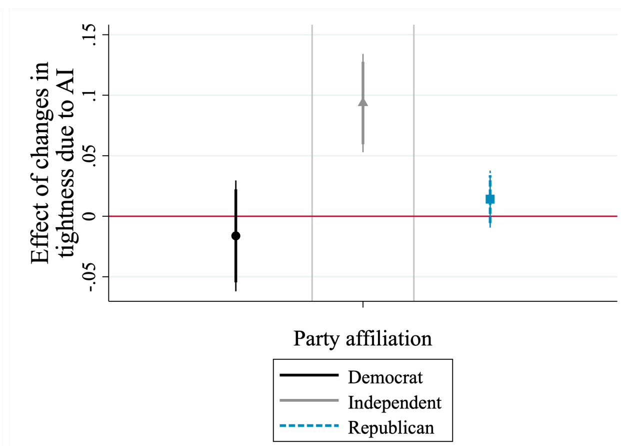

Notes: Confidence are clustered by county. * p<0.10, ** p<0.05, *** p<0.01. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

7 Conclusion

One of the fundamental global economic shifts of the last decade has been the emergence of scalable and commercially viable AI. We argue and show that AIs can benefit (incumbent) left-wing politicians if AI adoption benefits moderate voters, or if the demand for labor attracts more educated individuals—who are likely to vote for the left. We evaluate this premise using data from the

US' presidential elections, finding that there is an economic voting effect in favor of Democrats in Republican strongholds, which is driven by ideological moderates, and also evidence of a composition effect in favor of democrats in Democrat strongholds, which is explained by a higher share of highly educated individuals. Thus, we find evidence for both an economic voting and a composition effect in response to AI, electorally benefiting the left in the US.

As AI matures and further transforms labor markets across the world, the findings herein help us understand the political consequences of the future of work. Specifically, our findings indicate that while automation—like robot adoption—has bred support for the far-right in the past, new forms of automation like AI can and do breed support for the left. We show, however, that the timing and characteristics of automation matter. While the economic voting effect on moderates benefits the incumbent, which is coincidentally left-wing in our case study, the composition effect is more likely to benefit the left by increasing the share of educated individuals who are more likely to support liberal policies. In this regard, AI adoption has distinctive political effects vis-á-vis robot adoption as it affects different segments of labor markets and geographical areas. Hence, we highlight that different forms of automation can have different political effects.

It should be noted that our study captures the early deployment of AI, a stage at which firms are still developing the capacity to integrate these technologies into their business processes. As such, the political voting patterns we observe reflect only the initial phase of AI adoption. Going forward, however, these patterns are likely to evolve as the deployment of AI generates second-round effects that have not yet materialized. Recent studies of AI platforms, such as OpenAI (Chatterji et al., 2025) and Anthropic (Appel et al., 2025), document early usage patterns and foreshadow the types of tasks that may be subject to automation. While better-educated individuals and white-collar workers could face higher job-displacement risks over time, in the current stage, they still remain beneficiaries of early AI adoption. Since the debate on whether AI may cause widespread job displacement is still ongoing, as AI creates economies of scale and new opportunities for businesses and workers alike. Although AI has affected early career individuals, there are reasons to believe that these effects may be offset by its broad benefits in the near future (Mokyr, Vickers and Ziebarth, 2015; Autor, 2015; Aghion and Bunel, 2024).

References

Abadie, Alberto, Susan Athey, Guido W Imbens and Jeffrey M Wooldridge. 2023. "When should you adjust standard errors for clustering?" *The Quarterly Journal of Economics* 138(1):1–35.

Abraham, Katharine G, John C Haltiwanger and Lea E Rendell. 2020. "How tight is the US labor market?" *Brookings Papers on Economic Activity* 2020(1):97–165.

Acemoglu, Daron. 2025. "The simple macroeconomics of AI." *Economic Policy* 40(121):13–58.

Acemoglu, Daron, David Autor, Jonathon Hazell and Pascual Restrepo. 2022. "Artificial intelligence and jobs: Evidence from online vacancies." *Journal of Labor Economics* 40(S1):S293–S340.

Acemoglu, Daron, Gary W Anderson, David N Beede, Cathy Buffington, Eric E Childress, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, Pascual Restrepo and Nikolas Zolas. 2022. "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey.".

Acemoglu, Daron and Pascual Restrepo. 2020. "Robots and jobs: Evidence from US labor markets." *Journal of Political Economy* 128(6):2188–2244.

Aghion, Philippe, Céline Antonin, Simon Bunel and Xavier Jaravel. 2022. "The effects of automation on labor demand." *Robots and AI: A new economic era* pp. 15–39.

Aghion, Philippe and Peter Howitt. 2023. "The creative destruction approach to growth economics." *European Review* 31(4):312–325.

Aghion, Philippe and Simon Bunel. 2024. "AI and Growth: where do we stand?" *Policy Note* .

Agrawal Ajay, Joshua Gans, Avi Goldfarb. 2019. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Andreadis, Lefteris, Eleni Kalotychou, Manolis Chatzikonstantinou, Christodoulos Louca and Christos A Makridis. 2025. Local Heterogeneity in Artificial Intelligence Jobs over Time and Space. In *AEA Papers and Proceedings*. Vol. 115 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 pp. 29–34.

Andreadis, Lefteris, Emmanouil Chatzikonstantinou, Elena Kalotychou, Christodoulos Louca and Christos Makridis. 2025. "The local effects of artificial intelligence labor investments: Evidence from the municipal bond market." *Forthcoming at American Economic Review P&P* .

Anelli, Massimo, Italo Colantone and Piero Stanig. 2018. We were the robots: Automation and voting behavior in western europe. Technical Report 2019-115 Bocconi University–Mimeo.

Antoniades, Alexis and Charles W Calomiris. 2020. "Mortgage market credit conditions and US Presidential elections." *European Journal of Political Economy* 64:101909.

Antoniades, Alexis, Manolis Chatzikonstantinou and Mohammed Shahmeer Ahman. 2024. "The convergence of occupations: Evidence from Online Job Posts.".

Antoniades, Alexis, Manolis Chatzikonstantinou and Olesia Savka. 2024. "Are A.I. jobs driving up the demand for stem education? Yes, but not for women.".

Appel, Ruth, Peter McCrory, Alex Tamkin, Miles McCain, Tyler Neylon and Michael Stern. 2025. The Anthropic Economic Index report: Uneven geographic and enterprise AI adoption. Technical report Anthropic.
URL: *https://www.anthropic.com/research/anthropic-economic-index-september-2025-report*

Arkhangelsky, Dmitry, Guido W Imbens, Lihua Lei and Xiaoman Luo. 2024. "Design-robust two-way-fixed-effects regression for panel data." *Quantitative Economics* 15(4):999–1034.

Autor, David H. 2015. "Why are there still so many jobs? The history and future of workplace automation." *Journal of economic perspectives* 29(3):3–30.

Babina, Tania, Anastassia Fedyk, Alex He and James Hodson. 2024. "Artificial intelligence, firm growth, and product innovation." *Journal of Financial Economics* 151:103745.

Baccini, Leonardo, Costin Ciobanu and Krysztof Pelc. 2024. "why different economics shocks have different policy effects." *Mimeo, McGill University* .

Baccini, Leonardo and Stephen Weymouth. 2021. "Gone for good: Deindustrialization, white voter backlash, and US presidential voting." *American Political Science Review* 115(2):550–567.

Bailey, Jack. 2019. "The fact remains: Party ID moderates how voters respond to economic change." *Electoral Studies* 61:102071.

Balcazar, Carlos Felipe. 2023. "Globalization, Unions and Robots: The Effects of Automation on the Power of Labor and Policymaking." *Available at SSRN 4574527* .

Balcazar., Carlos Felipe. 2025. "Distributive conflict in the age of AI: Theory and evidence from the advent of ChatGPT." *Mimeo, University College London* .

Bansak, Kirk, Jens Hainmueller, Daniel J Hopkins and Teppei Yamamoto. 2023. "Using conjoint experiments to analyze election outcomes: The essential role of the average marginal component effect." *Political Analysis* 31(4):500–518.

Barberis, Nicholas C. 2013. "Thirty years of prospect theory in economics: A review and assessment." *Journal of economic perspectives* 27(1):173–196.

Bareis, Jascha and Christian Katzenbach. 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.

Barnichon, Régis and Adam Hale Shapiro. 2024. "Phillips meets beveridge." *Journal of Monetary Economics* 148:103660.

Becher, Michael and Daniel Stegmuller. 2023. Global Competition, Local Unions, and Political Representation: Disentangling mechanisms. Technical report Unpublished manuscript.

Bernanke, Ben and Olivier Blanchard. 2025. “What caused the US pandemic-era inflation?” *American Economic Journal: Macroeconomics* 17(3):1–35.

Berriochoa, Kattalina M. 2025. “Partisanship and Place: Education Preferences Across the Urban–Rural Divide.” *Political Studies* p. 00323217251382020.

Boix, Carles. 2019. *Democratic Capitalism at the Crossroads: Technological Change and the Future of Politics*. Princeton University Press.

Borusyak, Kirill, Peter Hull and Xavier Jaravel. 2022. “Quasi-experimental shift-share research designs.” *The Review of Economic Studies* 89(1):181–213.

Borusyak, Kirill, Peter Hull and Xavier Jaravel. 2025. “A practical guide to shift-share instruments.” *Journal of Economic Perspectives* 39(1):181–204.

Borwein, S., B. Magistro, R.M. Alvarez, B. Bonikowski and P. Loewen. 2025. “The Potential for Political Backlash Against Generative AI.” *Forthcoming at Public Opinion Quarterly* .

Borwein, Sophie, Bart Bonikowski, Peter Loewen, Beatrice Magistro and Blake Lee-Whiting. 2024. “Who Can Assert Ownership Over Automation? Workplace Technological Change, Populist and Ethno-nationalist Rhetoric, and Candidate Support.” *Political Behavior* .

Brynjolfsson, Erik, Bharat Chandar and Ruyu Chen. 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 . . . .

Center, The Niehaus. 2022. The Backlash against Globalization: Whats Next? Technical report.

Chatterji, Aaron, Thomas Cunningham, David J Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan and Kevin Wadman. 2025. How People Use ChatGPT. Working Paper 34255 National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w34255*

Chaudoin, Stephen and Michael-David Mangini. 2022. “Robots, Foreigners, and Foreign Robots: Policy Responses to Automation and Trade.” Mimeo, Harvard University.

Chaudoin, Stephen and Michael-David Mangini. 2024. “Robots, Foreigners, and Foreign Robots: Policy Responses to Automation and Trade.”.

Cooper, Jasper, Sung Eun Kim and Johannes Urpelainen. 2018. “The broad impact of a narrow conflict: How natural resource windfalls shape policy and politics.” *The Journal of Politics* 80(2):630–646.

Dasgupta, Aditya and Elena Ramirez. 2025. “Explaining rural conservatism: political consequences of technological change in the Great Plains.” *American Political Science Review* 119(1):277–299.

DeLuca, Kevin, Daniel J Moskowitz and Benjamin Schneer. 2024. “A drag on the ticket? Estimating top-of-the-ticket effects on down-ballot races.” *American Journal of Political Science* .

Demirci, Ozge, Jonas Hannane and Xinrong Zhu. 2025. "Who is AI replacing? The impact of generative AI on online freelancing platforms." *Management Science* .

Eggers, Andrew C. 2014. "Partisanship and Electoral Accountability: Evidence from the UK Expenses Scandal." *Quarterly Journal of Political Science* 9(4):441–472.
URL: *https://doi.org/10.1561/100.00013140*

Erikson, Robert S. 2009. "Economic Conditions and the Presidential Vote." *PS: Political Science & Politics* 42(3):467–472.

Fang, Tommy Pan and Shane Greenstein. 2025. Where the Cloud Rests: The Economic Geography of Data Centers. Working Paper 21-042 Harvard Business School. Revised August 2025.
URL: *https://www.hbs.edu/faculty/Pages/item.aspx?num=59368*

Fast, Ethan and Eric Horvitz. 2017. "Long-Term Trends in the Public Perception of Artificial Intelligence." Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence.
URL: *www.aaai.org*

Fedaseyeu, Viktar, Erik Gilje and Philip E. Strahan. 2015. Voter Preferences and Political Change: Evidence from Shale Booms. NBER Working Paper 21789 National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w21789*

Finseraas, Henning and Ole Henning Nyhus. 2025. "The political consequences of technological change that benefits low-skilled workers." *Political Science Research and Methods* 13(2):298–314.

Frank, Morgan R, David Autor, James E Bessen, Erik Brynjolfsson, Manuel Cebrian et al. 2019. "Toward understanding the impact of artificial intelligence on labor." *Proceedings of the National Academy of Sciences* 116(14):6531–6539.

Gallego, Aina, Alexander Kuo, Dulce Manzano and José Fernández-Albertos. 2021. "Technological Risk and Policy Preferences." *Comparative Political Studies :* pp. 1–33.

Gallego, Aina and Thomas Kurer. 2022. "Automation, Digitalization, and Artificial Intelligence in the Workplace: Implications for Political Behavior." *Annual Review of Political Science* 25:463–484.

Gazmararian, Alexander F. 2025. "Sources of partisan change: Evidence from the shale gas shock in American coal country." *The Journal of Politics* 87(2):601–615.

Gerber, Alan S, Gregory A Huber and Ebonya Washington. 2010. "Party affiliation, partisanship, and political beliefs: A field experiment." *American Political Science Review* 104(4):720–744.

Gimbel, Martha, Molly Kinder, Joshua Kendall and Maddie Lee. 2025. Evaluating the impact of AI on the labor market: Current state of affairs. Technical report Technical report, The Budget Lab at Yale University. URL: https://budgetlab . . . .

Goldsmith-Pinkham, Paul, Isaac Sorkin and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review* .

Guntermann, Eric, Gabriel S Lenz and Jeffrey R Myers. 2021. "The impact of the economy on presidential elections throughout US history." *Political Behavior* 43(2):837–857.

Hainmueller, Jens and Michael J Hiscox. 2006. "Learning to love globalization: Education and individual attitudes toward international trade." *International Organization* 60(2):469–498.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence DW Schmidt and Bryan Seegmiller. 2025. Artificial intelligence and the labor market. Technical report National Bureau of Economic Research.

Healy, Andrew and Neil Malhotra. 2013. "Retrospective voting reconsidered." *Annual review of political science* 16(1):285–306.

Holliday, Derek. 2022. "Is all politics presidential? Decomposing partisan patterns in US Election outcomes across offices, 1972–2020.".

Imbens, Guido W. 2003. "Sensitivity to exogeneity assumptions in program evaluation." *American Economic Review* 93(2):126–132.

Inglese, Gaetano. 2025. From Aspiration to Disillusion? The Political Consequences of AI-driven Employment Threats.

Jaeger, David A, Joakim Ruist and Jan Stuhler. 2018. Shift-share instruments and the impact of immigration. Technical report National Bureau of Economic Research.

Jensen, J. Bradford, Dennis P. Quinn and Stephen Weymouth. 2017. "Winners and Losers in International Trade: The Effects on US Presidential Voting." *International Organization* 71(3):423–457.

Jones, Charles I. 2024. "The ai dilemma: Growth versus existential risk." *American Economic Review: Insights* 6(4):575–590.

Kayser, Mark Andreas and Christopher Wlezien. 2011. "Performance Pressure: Patterns of Partisanship and the Economic Vote." *European Journal of Political Research* 50(3):365–394.
URL: *https://doi.org/10.1111/j.1475-6765.2010.01937.x*

Kelley, Patrick Gage, Yongwei Yang, Courtney Heldreth, Christopher Moessner, Aaron Sedley, Andreas Kramm, David T. Newman and Allison Woodruff. 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*. Association for Computing Machinery, Inc pp. 627–637.

Lepers, Etienne. 2023. "Surfing the credit wave: Government popularity as driver of credit cycles." *European Journal of Political Economy* p. 102456.

Lewis-Beck, Michael S. and Charles Tien. 2008. "The Job of President and the Jobs Model Forecast: Obama for '08?" *PS: Political Science & Politics* 41(4):687–690.

Magistro, Beatrice, Peter Loewen, Bart Bonikowski, Sophie Borwein and Blake Lee-Whiting. 2024. "Attitudes toward automation and the demand for policies addressing job loss: the effects of information about trade-offs." *Political Science Research and Methods* .

Magistro, Beatrice, Sophie Borwein, R Michael Alvarez, Bart Bonikowski and Peter John Loewen. 2025. "Attitudes toward artificial intelligence (AI) and globalization: Common microfoundations and political implications." *American Journal of Political Science* .

McElheran, Kristina, J Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster and Nikolas Zolas. 2024. "AI adoption in America: Who, what, and where." *Journal of Economics & Management Strategy* 33(2):375–415.

Mellon, Jonathan. 2013. "Where and when can we use Google Trends to measure issue salience?" *PS: Political Science & Politics* 46(2):280–290.

Menon, Anil and Baobao Zhang. 2024. "Messaging about different AI development timelines fails to change the public's policy preferences.".

Menon, Anil and Baobao Zhang. 2025. "Future Shock or Future Shrug? Public Responses to Varied Artificial Intelligence Development Timelines.".

Merz, Nicolas, Sven Regel and Jirka Lewandowski. 2016. "The Manifesto Corpus: A new resource for research on political parties and quantitative text analysis." *Research & Politics* 3(2):2053168016643346.

Mettler, Suzanne and Trevor Brown. 2022. "The growing rural-urban political divide and democratic vulnerability." *The ANNALS of the American Academy of Political and Social Science* 699(1):130–142.

Milanez, Anna. 2023. OECD Social, Employment And Migration Working Papers No. 288.
URL: *http://dx.doi.org/10.1787/2247ce58-en*

Mokyr, Joel, Chris Vickers and Nicolas L Ziebarth. 2015. "The history of technological anxiety and the future of economic growth: Is this time different?" *Journal of Economic Perspectives* 29(3):31–50.

Naoi, Megumi. 2020. "Survey experiments in international political economy: what we (don't) know about the backlash against globalization." *Annual Review of Political Science* 23(1):333–356.

Neri, Hugo and Fabio Cozman. 2020. "The role of experts in the public perception of risk of artificial intelligence." *AI and Society* 35:663–673.

Pissarides, Christopher A. 2011. "Equilibrium in the labor market with search frictions." *American Economic Review* 101(4):1092–1105.

Raviv, Shir. 2024. "How Media Coverage and Elite Communication Shape Public Opinion on AI Regulation * Tamar Mitts.".
URL: *https://doi.org/10.17605/OSF.IO/UQDXB*

Reauscher, Natalie. 2022. *The Future of Work in the United States: Discourses on Automation and the Platform Economy*. Springer.
URL: *http://www.springer.com/series/1262*

Roe, Jasper and Mike Perkins. 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.

Ryazanov, Igor, Carl Öhman and Johanna Björklund. 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.

Scott, Ralph. 2022. "Does university make you more liberal? Estimating the within-individual effects of higher education on political values." *Electoral Studies* 77:102471.

Wu, Nicole. 2018. "Misattributed blame? Attitudes toward globalization in the age of automation." *Political Science Research and Methods* pp. 1–18.

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

Wu, Nicole. 2024. "Foreign competition and technological attitudes in the United States *.".

Zhang, Baobao. 2022. No rage against the machines: Threat of automation does not change policy preferences. In *AIES 2022 - Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, Inc pp. 856–866.

Zhang, Baobao. 2023. Public Opinion toward Artificial Intelligence. In *The Oxford Handbook of AI Governance*. Oxford University Press.
URL: *https://doi.org/10.1093/oxfordhb/9780197579329.013.36*

Zingher, Joshua N. 2022. "TRENDS: Diploma divide: Educational attainment and the realignment of the American electorate." *Political Research Quarterly* 75(2):263–277.

.tocmtappendix mtchapternone mtappendixsubsection mtappendixsubsubsection

Online Appendix

Table of Contents

A Exogeneity and treatment response

While AI's development took place over decades, scholars observed that the release of ChatGPT was the milestone that catapulted generative AI into the mainstream. First, we use data from private sector lay-offs from the Bureau of Labor statistics to show that there was a substantial increase in lay-offs around the time where ChatGPT is released (november of 2022), as opposed to in 2017 when the transformer architecture was released (Figure A1); the sharp increase in lay-offs about 2019-2020 is due to the Covid19 pandemic.

Figure A1: Private sector lay-offs

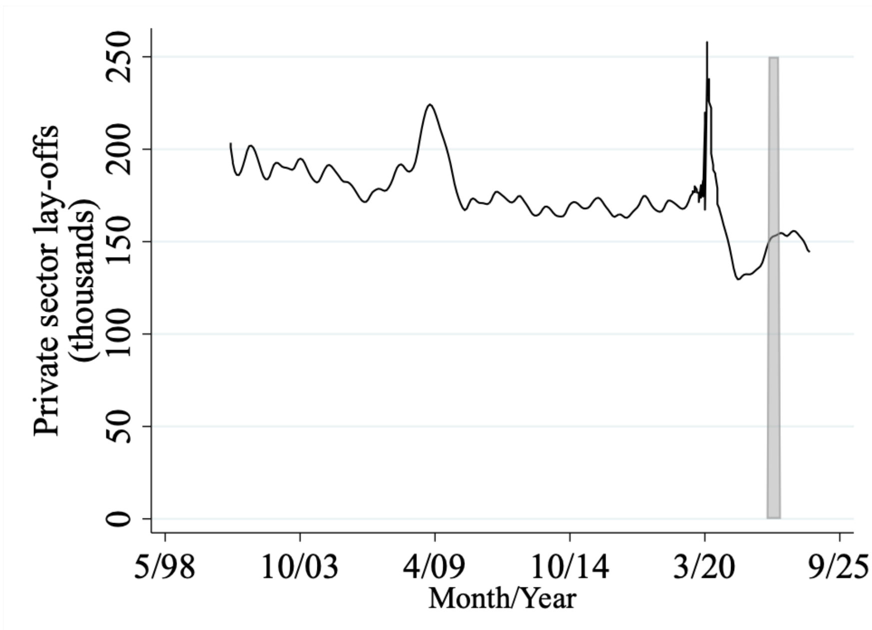

Furthermore, following Balcazar. (2025) we use Google trends to measure the extent of people's interest (excitement, worry, etc.) in regards to AI technologies, which is a strategy that has been used to measure interest in other areas (Mellon, 2013). Additionally, we measure related news supply, using the Global Database of Events, Language and Tone (GDELT), which is real-time database of newsworthy events.

We observe widespread interest and news coverage regarding GPT starting at the release of ChatGPT, despite the transformer structure became available in 2017. GPTs were not commercial at a big scale until then; only niche occupations such as researchers in computer science and other tech workers, especially in silicon valley, were aware of these technological developments. Figure A2, panel a, gives credence to the quasi-exogenous nature of the event. In it we use 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, an later on around the release of GPT4 on March 14 of 2023. Figure A2, panel b, shows the news interest regarding the keywords above We observe that news coverage about GPTs is lagged vis-á-vis interest from people using Google trends. However after the release of GPT4, the news coverage switches to using AI as a term.

Figure A2: Normalized Google trends interest scores for ChatGPT and AI

(a) Google trends **(b) GDELT**

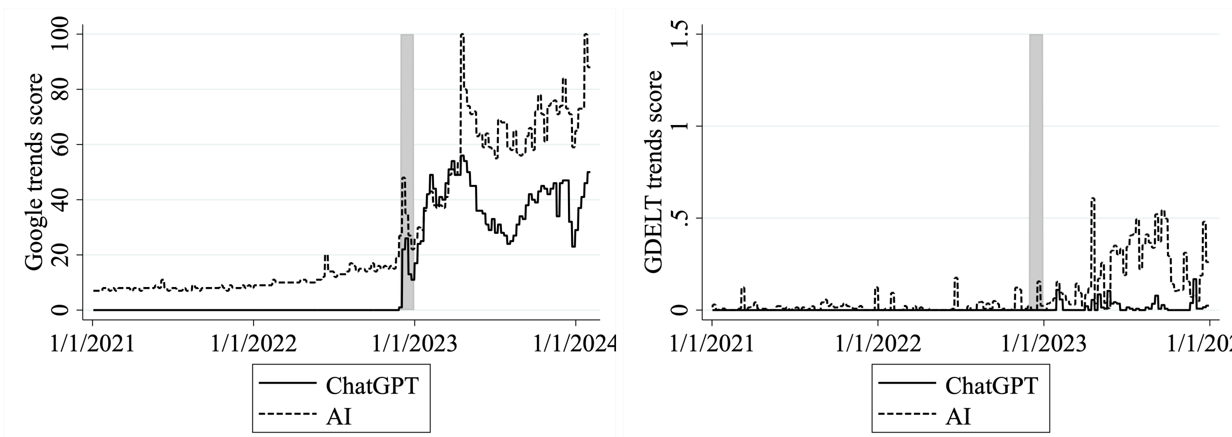

Note: Google trends interest scores for each country in the world are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. GDELT uses data from the Internet Archive's Television News Archive. I focus on data from networks with international coverage such as CNN, Al Jazeera, BBC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

We also observe a correlation between our measure of the AI shock and google search queries for the period 2016-2024 (Figure A3). While this suggests a correlation between people's interest in AI and the AI shock, we do not find evidence for a similar residualized correlation.

Lastly, we evaluate the possibility that other shocks such as robot adoption, offshoring or chinese import competition could also be associated with AI intensity. We don't expect this to be the case because the spatial distribution of AI vacancies is not centered in previously manufacturing areas like the Rust Belt. Second we find a negative correlation between these shocks and our mea-

Figure A3: Correlations between changes in AI intensity and changes in google search queries

(a) Not residualized **(b) Residualized**

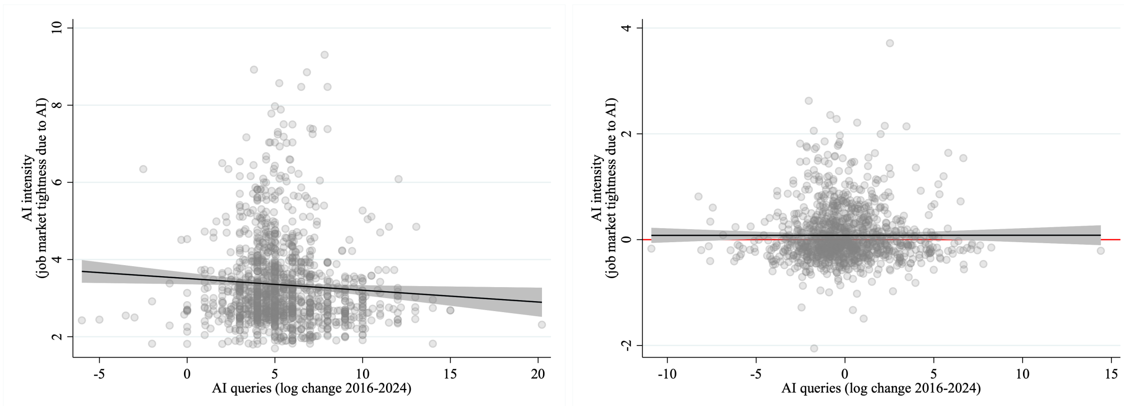

Note: Google trends interest scores for each country in the world are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. GDELT uses data from the Internet Archive's Television News Archive. I focus on data from networks with international coverage such as CNN, Al Jazeera, BBC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

sure of AI intensity, although these correlations are not robust to residualization (Figure A4). This is confirmed in our sensitivity analysis below.

Figure A4: Correlations between AI shock and other shocks

(a) Robot adoption
(not residualized)

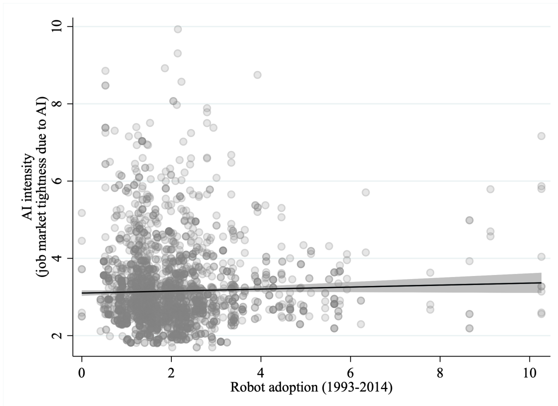

(b) Robot adoption
(residualized)

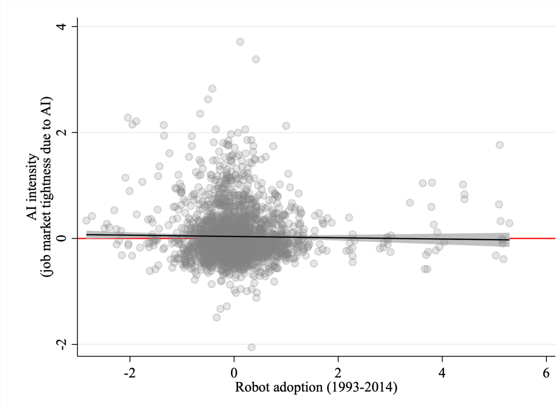

(c) Offshoring
(not residualized)

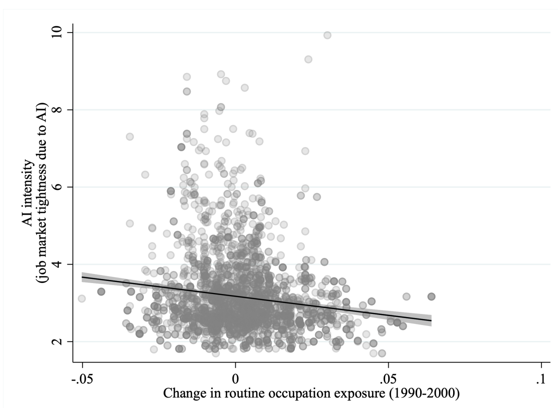

(d) Offshoring
(residualized)

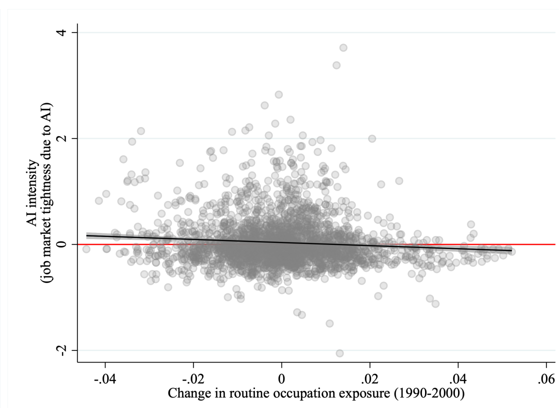

(e) Import competition
(not residualized)

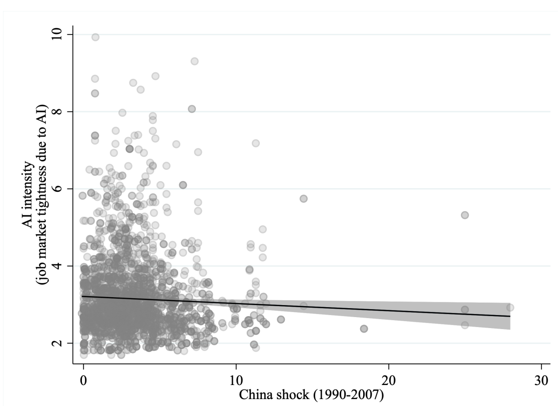

(f) Import competition
(residualized)

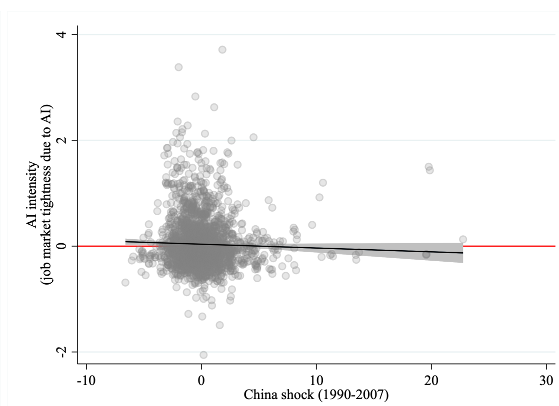

Note: Google trends interest scores for each county in are obtained from using ChatGPT as a keyword between 01/01/2021 and 01/31/2024. All other shocks are measured for theior respective time periods in order to evaluate long-run secular trends.

B Robustness

Three relevant concerns may arise from the empirical design: i) The possibility that confounders that are not accounted for may overturn the results, ii) The possibility that any single observation is driving the results, iv) Sensitivity to measurement, iv) The possibility of spurious correlations, v) The possibility of a composition effect. We explore these potential issues next:

Sensitivity to unobserved confounding. We 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 to treatment and the outcome variable simultaneously that the estimated coefficients may not be robust to. We explicitly relax the exogeneity assumption by allowing for a limited amount of correlation between treatment and unobserved components of the outcome (Imbens, 2003). We find that an unobservable confounder that could potentially overturn the main results needs to exhibit a higher partial R^2 vis-á-vis than the confounders already included (Figure B1). This is unlikely to exist since it would need to have a much stronger effect than import competition— the confounder with the highest partial R^2.

Parameter stability to observations. To further corroborate that there's no one observation driving the results, we carry out a robustness tests wherein we drop on congressional district at a time with replacement (*á la* Jackknife). We find that the treatment indicator is quite stable and statistically significant for each permutation (Figure B2).

Alternative measures of AI intensity. Computer-related job vacancies exhibit the highest concentration of AI skills (Antoniades, Chatzikonstantinou and Ahman, 2024). Thus now we use the broader definition that may encompass general IT roles unrelated to AI (Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo and Zolas, 2022). We find similar results for the alternative measure (Figure B3).

2016 and 2020 elections. We have claimed that is the advent of ChatGPT and not the transformer architecture the driver of the AI adoption shock. As such we shouldn't observe consistent

Figure B1: Sensitivity analysis to unobserved confounding (adding shocks)

(a) Vote share for democrats

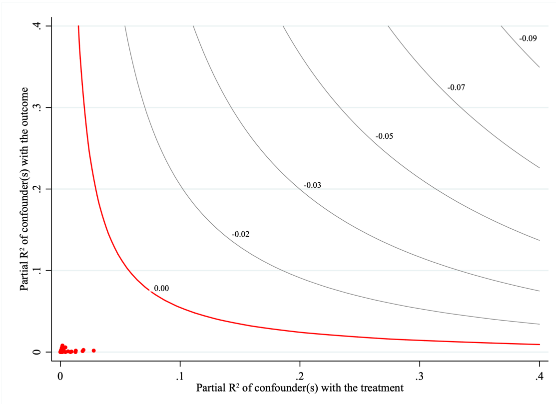

(b) Democrats flip county

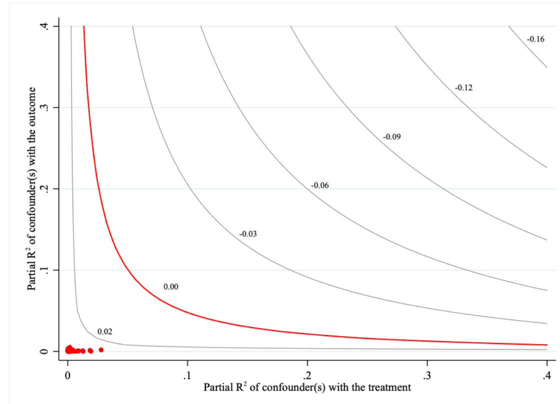

(c) Close race

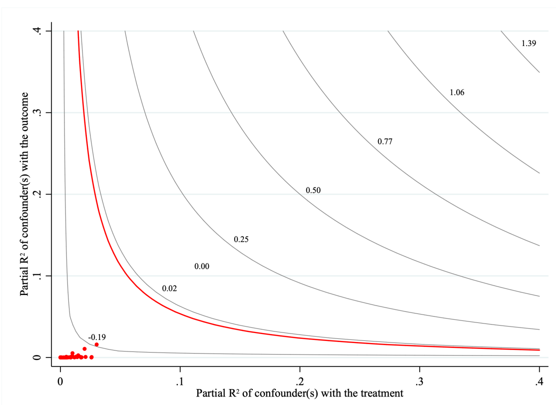

(d) Margin of victory

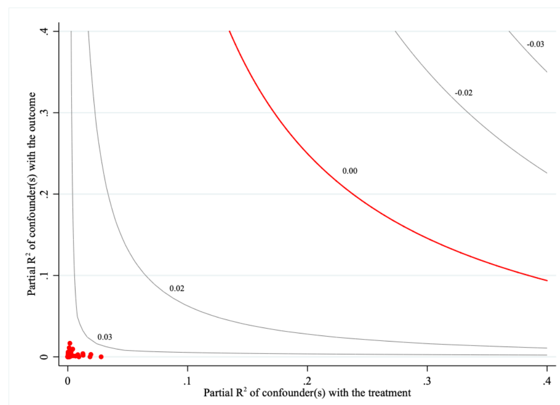

Standard errors clustered by district in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location × time non-monotonic trends given that our design is in the first differences.

and statistically significant effects if we use the 2016 and 2020 elections. Table B1 shows the results of this placebo test, confirming it.

Composition effects. The increased demand for AI could drive our results, if this attracts more liberal workers into previously redder areas. To explore this possibility further, we check whether changing socio-demographics of the population, toward minorities and highly educated individuals—which are more likely to support democrats. Figure B4 displays our results. It shows

Figure B2: Parameter stability to excluding one district with replacement

(a) Vote share for democrats

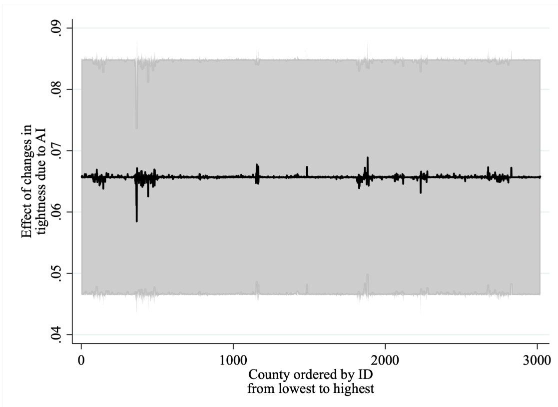

(b) Democrats flip county

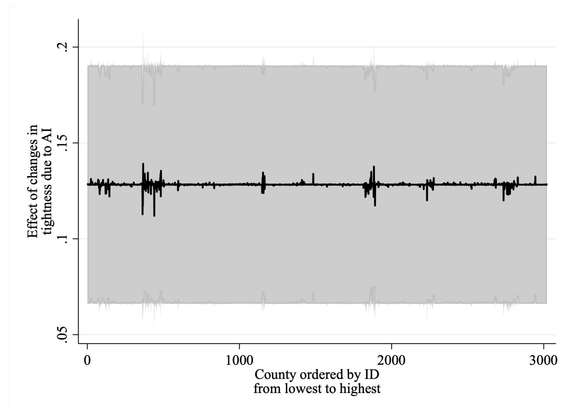

(c) Close race

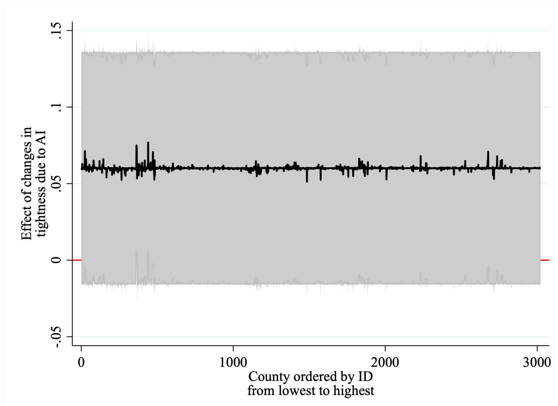

(d) Margin of victory

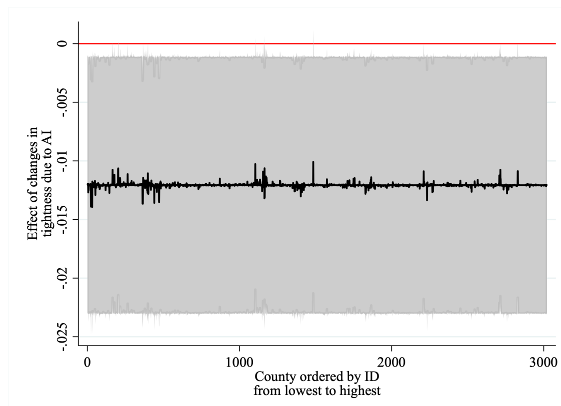

Confidence bands are clustered by district. * p<0.10, ** p<0.05, *** p<0.01. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

that there is little evifence for other composition effects.

We also check that our results are not driven by the distribution of the rural population by controlling by the interaction between the share of those employed in agriculture in 2010, which tend to lean more Republican, and our treatment (Figure B5). In other words, we demonstrate absence of a bundled mechanism.

Using panel data. We use the recently published panel from the CCES to further explore composition effects. We focus on those individuals that can be followed between 2020 and 2024—

Figure B3: Heterogeneous Effects of AI by revealed share of democrats, including software-related jobs

(a) Democrat vote share

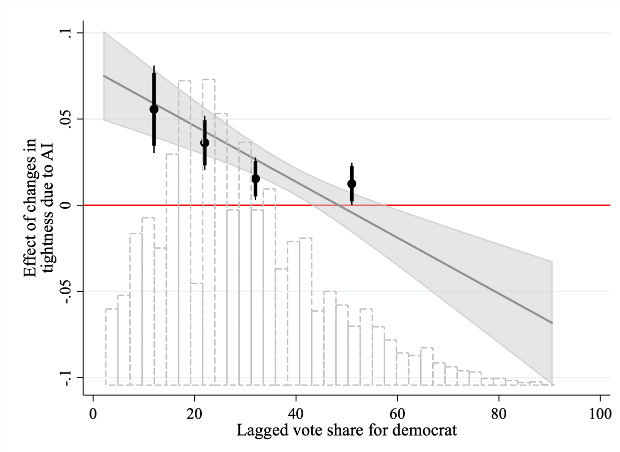

(b) Flips county

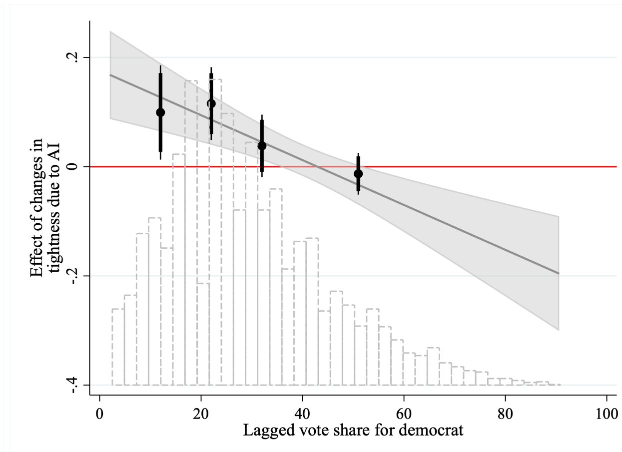

(c) Close race (<=0.05 margin)

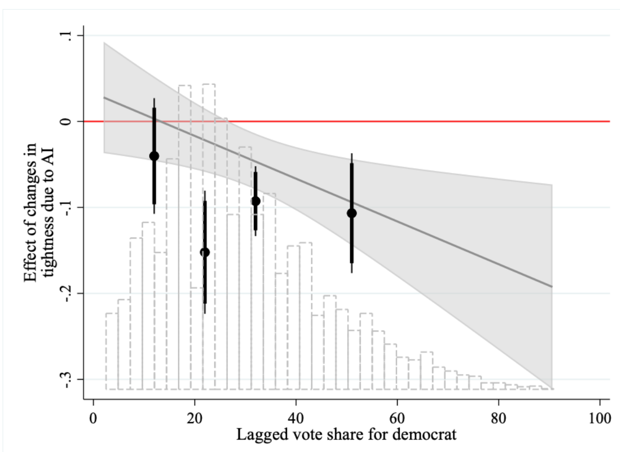

(d) Democrat margin of victory

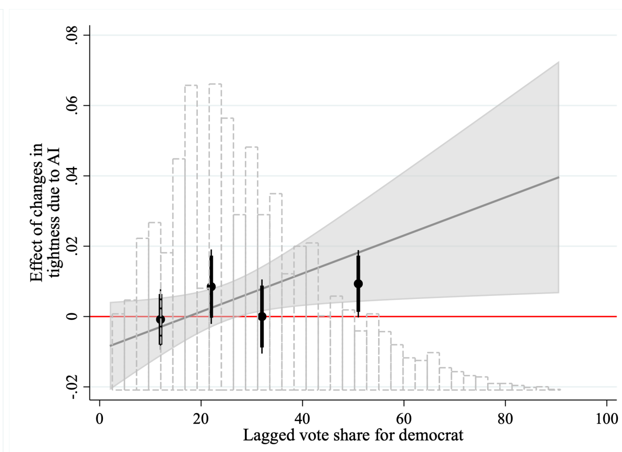

Notes: Confidence intervals and confidence bands are clustered by district. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

about six thousand individuals. For these individuals we know if they report living in different congressional districts at the time of the survey. We estimate the basic profiles of the individuals that moved to areas that saw an increase in job market intensity due to AI. To do this we assign each individual the AI shock of the area they moved into, as well as all other socio-demographics and socio-economic pre-treatment variables. Then we estimate the moderating effect of the AI

Table B1: Effect of AI intensity on presidential outcomes (2016-2020)

	Outcomes for democrats			
	Vote share	*Flips county*	*Close race*	*Margin of Victory*
	(1)	(2)	(3)	(4)
AI intensity	0.0041	0.0271**	-0.0143	-0.0004
	(0.0028)	(0.0129)	(0.0110)	(0.0019)
Observations	2747	2747	2747	2747
Congressional District FE	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes
Industry controls	Yes	Yes	Yes	Yes
Other shocks	Yes	Yes	Yes	Yes
Union power	Yes	Yes	Yes	Yes

Notes: Standard errors clustered by district in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

shock conditional on mover status, on education, ideology and party identification.

We show the results of our estimations in Table B2. We observe that those that moved to areas with more AI adoption were highly educated individuals, which are on average more likely to support liberal policies, and thus liberal candidates (Hainmueller and Hiscox, 2006; Merz, Regel and Lewandowski, 2016; Zingher, 2022).

Table B2: Profiles of those who moved congressional districts

	Outcomes for individuals		
	Has higher education	***Is liberal***	***Is democrat***
	(1)	(2)	(3)
AI shock	0.0102 (0.0334)	0.0345 (0.0499)	0.1069** (0.0485)
AI shock x moved district	0.0926*** (0.0354)	-0.0002 (0.0508)	-0.1057** (0.0522)
Observations	5922	5922	5922 [b]
State District FE	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes
Industry controls	Yes	Yes	Yes
Other shocks	Yes	Yes	Yes
Union power	Yes	Yes	Yes

Notes: Standard errors clustered by district in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

Figure B4: Heterogeneous Effects of AI by revealed share of democrats on composition of the population

(a) Change in share of whites

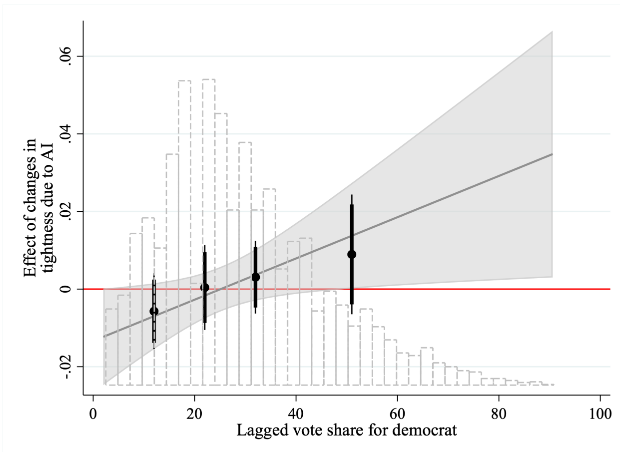

(b) Change in share of people with only high school education

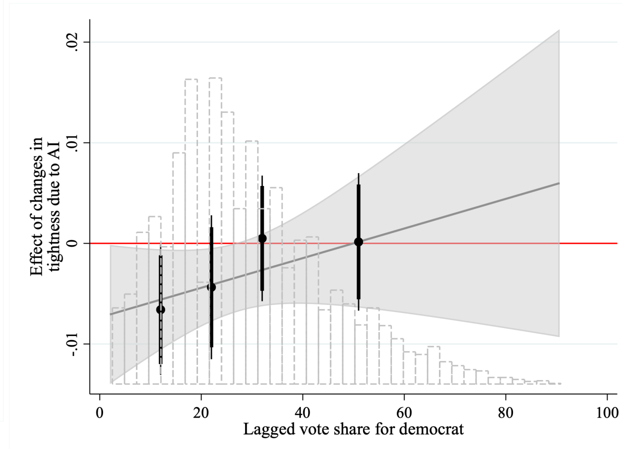

(c) Change in median age

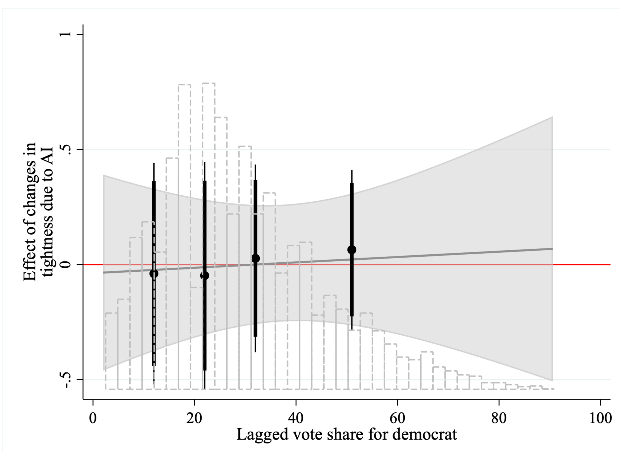

(d) Change in share of people 25 years or more

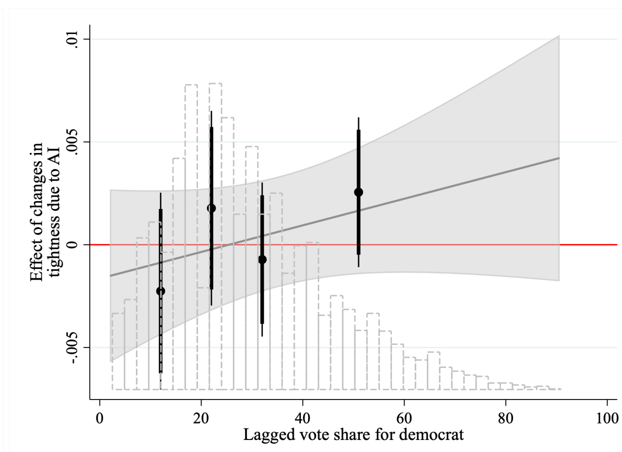

(e) Change in median house value

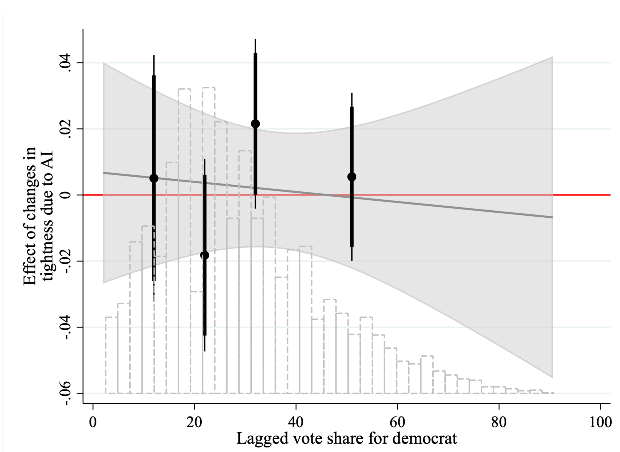

(f) Change in median rent

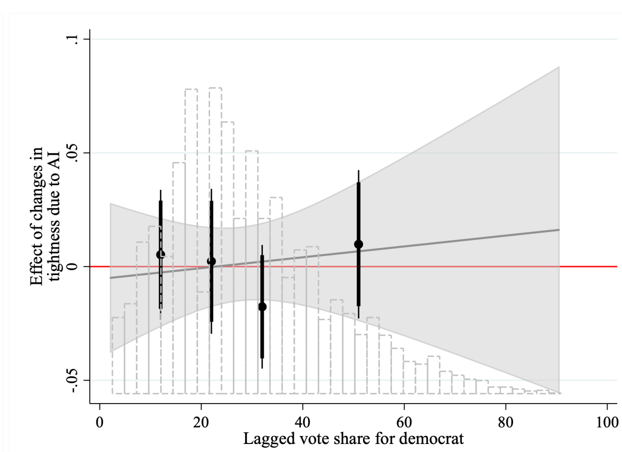

Notes: Confidence intervals and confidence bands are clustered by district. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

Figure B5: Heterogeneous Effects of AI intensity on approval for president by revealed share of democrats, ruling out a rural population composition mechanism

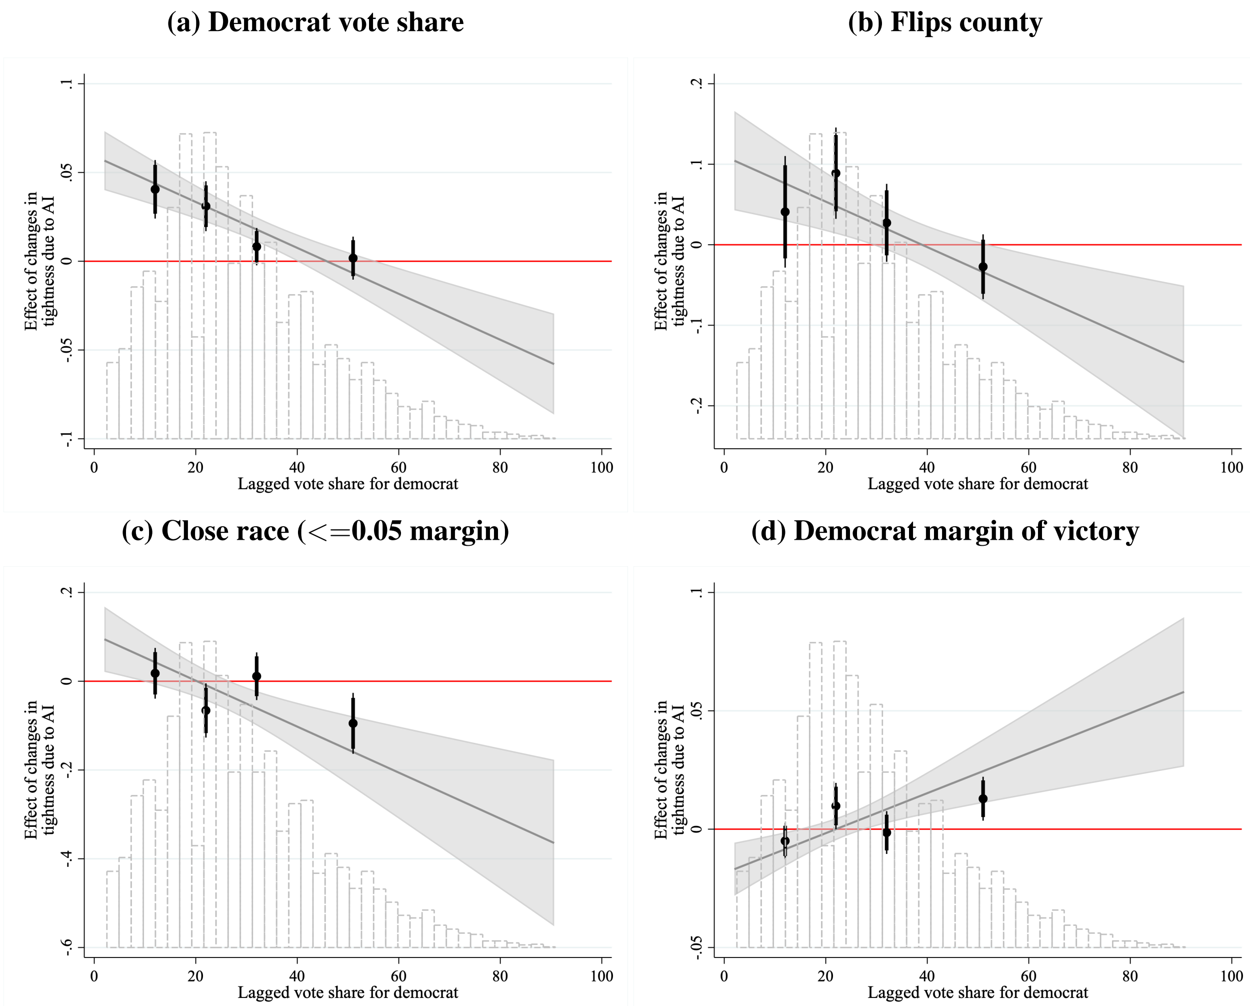

Notes: Confidence intervals and confidence bands are clustered by district. Estimates are heterogeneity robust. Controls include changes in both the socio-economic and socio-demographic composition at the county level, pre-treatment changes in unionization, and other economic shocks such as the the China import shock, the robot adoption shock, deindustrialization, and Covid 19. We also control for multiple secular trends by controlling by congressional-district fixed effects, which capture location $\times$ time non-monotonic trends given that our design is in the first differences.

C Descriptive statistics

Table C1: Summary statistics

	All					AI shock				
	Average	**Sd.**	**Min**	**Max**	**Obs**	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**
Trteatment (as changes between 2019 and 2023):										
AI shock (in SDs)	2.12	1.00	0.70	10.12	3106.00	1.33	1.59	1.86	2.26	3.54
AI shock excluding IT occupations (in SDs)	3.14	1.00	1.38	9.81	3106.00	2.10	2.55	2.95	3.41	4.71
Outcomes (as changes between 2020 and 2024):										
Vote share (%)	5.03	12.43	-13.60	70.86	3106.00	5.07	4.85	4.62	3.17	7.45
Close race (<5% margin)	0.90	0.30	0.00	1.00	3106.00	0.95	0.95	0.95	0.91	0.72
Flipped county	0.04	0.20	0.00	1.00	3106.00	0.03	0.03	0.02	0.03	0.10
Margin of victory	1.64	5.22	0.00	43.27	3106.00	0.80	0.75	0.70	1.38	4.55
Lagged vote share for democrats (no changes)	27.08	16.88	2.21	92.15	3106.00	20.77	23.12	25.61	30.07	35.86
Composition (as changes between 2020 and 2024):										
Log population	0.00	0.07	-2.00	0.55	3106.00	-0.02	-0.01	0.00	0.01	0.02
Log median rent	0.16	0.09	-0.63	0.78	3089.00	0.16	0.16	0.16	0.17	0.18
Log median house value	0.26	0.10	-0.39	0.67	3086.00	0.24	0.26	0.26	0.27	0.27
Share of population ovetr 25 years (%)	0.27	1.52	-10.58	22.07	3106.00	0.12	0.24	0.32	0.26	0.42
Share with college degree or more (%)	1.48	2.14	-10.31	22.85	3106.00	1.15	1.50	1.40	1.54	1.79
Share with highschool degree (%)	1.17	1.99	-12.86	24.96	3106.00	1.70	1.33	1.12	0.94	0.76
Log median age	0.18	1.49	-11.50	24.30	3106.00	0.09	0.29	0.12	0.09	0.32
Covariates (as changes between 2018 and 2022):										
Log population	0.00	0.06	-1.20	0.31	3106.00	-0.02	-0.01	0.00	0.01	0.02
Share with college degree or more (%)	1.82	2.49	-9.08	22.76	3106.00	1.38	1.59	1.85	1.90	2.39
Share with highschool degree (%)	1.76	2.29	-22.90	15.40	3106.00	2.31	1.97	1.71	1.52	1.30
Share of foreign born workers (%)	0.04	1.24	-18.56	8.86	3106.00	-0.01	0.01	0.00	0.01	0.18
Share of black workers (%)	-0.09	1.12	-14.08	14.28	3106.00	-0.20	-0.18	-0.10	0.07	-0.05
Share of hispanic workers (%)	0.67	1.47	-18.45	15.31	3106.00	0.59	0.70	0.73	0.71	0.63
Share employed in civilian workers (%)	-0.15	2.40	-21.38	19.28	3106.00	-0.08	-0.15	-0.22	-0.29	-0.03
Share employed in agriculture (%)	-0.46	2.37	-22.88	20.12	3106.00	-0.57	-0.61	-0.29	-0.51	-0.31
Share employed in manufactures (%)	0.01	2.40	-17.89	21.06	3106.00	0.23	0.16	-0.17	-0.17	-0.01
Share of white workers (%)	-2.80	3.62	-36.94	24.66	3106.00	-2.47	-2.91	-2.88	-2.76	-3.00
Share of population ovetr 25 years (%)	0.58	1.67	-23.84	14.32	3106.00	0.45	0.60	0.48	0.55	0.82
Share of unemployed (%)	-0.51	0.95	-10.63	6.59	3106.00	-0.65	-0.58	-0.46	-0.42	-0.45

	All					Mean by quartiles of AI shock				
	Average	**Sd.**	**Min**	**Max**	**Obs**	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**
Covariates (as changes between 2018 and 2022):										
Import competition (China) shock	0.15	0.12	0.00	0.65	3039.00	0.16	0.15	0.13	0.13	0.17
Deindustrialization shock	0.23	0.17	0.00	0.85	3039.00	0.23	0.21	0.23	0.22	0.26
Manufacturing lay-offs (%)	-7.33	8.05	-36.24	45.18	3039.00	-8.82	-6.67	-6.96	-7.03	-7.13
Interest in AI (G-search score)	67.14	3.66	36.38	75.11	2983.00	66.87	67.14	67.34	67.18	67.16
Robot adoption shock	0.33	0.30	-0.01	2.08	3085.00	0.37	0.31	0.34	0.31	0.31
Union members (thou.)	-0.09	1.03	-1.36	14.12	3106.00	-0.01	-0.01	-0.02	-0.04	-0.35
Covid 19 deaths (thou.)	0.03	0.10	0.00	3.47	3101.00	0.02	0.02	0.02	0.03	0.08
Covid 19 cases (thou.)	3.05	10.58	0.00	363.24	3101.00	1.03	1.45	1.74	2.90	8.11
No. agriculture businesses (thou.)	0.00	0.03	-0.24	0.36	3023.00	0.00	0.00	0.00	0.01	0.01
No. ICT businesses (thou.)	0.02	0.09	-0.96	3.19	3023.00	0.00	0.01	0.01	0.01	0.04
No. finance businesses (thou.)	0.01	0.13	-1.63	4.90	3023.00	0.00	0.01	0.01	0.02	0.02
No. science businesses (thou.)	0.11	0.59	-4.77	15.95	3023.00	0.02	0.07	0.07	0.14	0.27
Payroll agriculture (mill.)	0.64	3.58	-26.45	85.31	3023.00	0.40	0.44	0.47	0.71	1.17
Payroll ICTs (mill.)	47.15	824.50	-340.45	24512.12	3023.00	0.08	0.47	0.47	6.68	226.03
Payroll finance (mill.)	51.18	421.73	-277.59	17257.19	3023.00	1.89	5.83	7.09	32.62	206.78
Payroll science (mill.)	74.16	585.18	-122.91	15613.80	3023.00	1.91	5.56	7.11	28.14	325.29

Table C2: Profiles by ideology

	All						Ideology	[t]
	Average	**Sd.**	**Min**	**Max**	**N**	**Liberal**	**Moderate**	**Conservative** [b]
								[t]
Age	50.43	17.51	16.00	96.00	120000.00	47.82	48.78	55.44
Gender	-0.51	0.56	-1.00	3.00	120000.00	-0.52	-0.54	-0.47
Race								
White	0.69	0.46	0.00	1.00	120000.00	0.68	0.62	0.80
Black	0.13	0.34	0.00	1.00	120000.00	0.14	0.17	0.06
Hispanice	0.09	0.28	0.00	1.00	120000.00	0.09	0.10	0.06
Asian	0.03	0.17	0.00	1.00	120000.00	0.04	0.03	0.02
Other	0.06	0.24	0.00	1.00	120000.00	0.06	0.07	0.06
Education								
Incomplete secondary edu.	0.04	0.20	0.00	1.00	120000.00	0.02	0.06	0.03
Secondary edu.	0.50	0.50	0.00	1.00	120000.00	0.42	0.53	0.55
Higher edu.	0.33	0.47	0.00	1.00	120000.00	0.38	0.30	0.32
Postgraduate	0.13	0.33	0.00	1.00	120000.00	0.18	0.11	0.10
Employment								
Union Membership	0.25	0.43	0.00	1.00	119910.00	0.27	0.22	0.24
Unemployed	0.08	0.27	0.00	1.00	119774.00	0.07	0.10	0.05
Employed	0.48	0.50	0.00	1.00	119774.00	0.52	0.48	0.45
Other	0.44	0.50	0.00	1.00	119774.00	0.41	0.42	0.50
Party identification								
Democrat	0.39	0.49	0.00	1.00	120000.00	0.76	0.32	0.07
Independent	0.36	0.48	0.00	1.00	120000.00	0.22	0.54	0.26
Republican	0.26	0.44	0.00	1.00	120000.00	0.02	0.14	0.67
Political opinion								
Approves the President	0.20	0.40	0.00	1.00	116490.00	0.38	0.17	0.04
Economy is doing better	0.23	0.42	0.00	1.00	116528.00	0.43	0.21	0.06
Supports immigration	0.66	0.47	0.00	1.00	119961.00	0.90	0.69	0.37
Supports abortion rights	0.64	0.48	0.00	1.00	119929.00	0.93	0.70	0.27
Racism is a problem	0.60	0.49	0.00	1.00	98918.00	0.89	0.61	0.31 [b]

D Survey Questionnaires

Table D1: CCES survey questions

Question	Options	Variable name	Value
Would you say that OVER THE PAST YEAR the nation's economy has	Gotten much better	Approves the president	1
	Gotten somewhat better		1
	Stayed about the same		0
	Gotten somewhat worse		0
	Gotten much worse		0
Do you approve of the way each is doing their job? (President Biden)	Strongly approve	State of the economy	1
	Somewhat approve		1
	Somewhat disapprove		0
	Strongly disapprove		0
Congress has been debating different policiess concerning immigration reform. The Senate proposal has a path to citizenship for illegal immigrants.	Stricter migration	Support immigration	0
	Only for legal migrants		1
There has been some discussion about abortion during recent years. Which one of the opinions on this page best agrees with your view on this issue?	No	Support abortion	0
	For rape, incest or threat to life		0
	When need is established		0
	Yes		1
Racial problems in the U.S. are rare, isolated situations.	Strongly agree	Racism is a problem	1
	Somewhat agree		1
	Somewhat disagree		0
	Strongly disagree		0

E AI-Related Skills

Table E1: AI-Related Lightcast Skills

Apache Ant	Convolutional Neural Network (CNN)
Apache Spark	Decision Trees
Artificial Intelligence	Deep Learning
Automated Testing	Deeplearning4j
Automation Consulting	Machine Learning
Automation Systems	Machine Vision
Automation Techniques	MapReduce
Automation Tools	Kubernetes
AWS Elastic MapReduce (EMR)	Natural Language Processing
BigQuery	Natural Language Toolkit (NLTK)
Boosting (Machine Learning)	Neural Networks
Caffe Deep Learning Framework	Splunk
Cluster Analysis	Supervised Learning (Machine Learning)
Clustering	Support Vector Machines (SVM)
Clustering Algorithms	TensorFlow
Computer Vision	Torch (Machine Learning)