UC Davis

UC Davis Electronic Theses and Dissertations

Title

Economic Shocks, Automation Threats, and the New Politics of The Welfare State in the AI-era

Permalink

https://escholarship.org/uc/item/7jd552th

ISBN

9798193261940

Author

Haddad, Julian Anthony

Publication Date

2026-09-08

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Economic Shocks, Automation Threats, and the New Politics of The Welfare State in the AI-era

By

JULIAN HADDAD

DISSERTATION
Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Political Science

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS
Approved:

Daniel Kono, Chair

Lauren Young

Hanno Hilbig

Christopher Witko

Committee in Charge

2026

Abstract

In 2022 the OECD reported that member countries spent an average of 21% of GDP on public social spending. Indeed, welfare states make up a massive portion of national budgets and serve as an important safety net in times of economic crisis. Unsurprisingly, welfare spending is a bitterly fought-over political arena. As the political reverberations of the Global Financial Crisis continue to ripple and artificial intelligence heralds another potentially destabilizing economic shock, understanding the politics of the welfare state in the 21st century is paramount. This dissertation offers contributions to political science in two main areas: first, to the literature that argues existing welfare state contexts are a key variable in understanding political behavior in response to economic shocks; and second, to research on attitudes toward welfare state expansion in the face of automation threats. Chapter 1 offers causal evidence to answer how attitudes toward the welfare state and the regulation of automation technologies are shaped by AI job threats. I conducted a survey experiment (n=1,009) randomly showing participants videos with information about AI as either a productivity tool or as a job threat. I find that policy attitudes shift along partisan and educational lines. In response to all AI information, Republicans prefer to reduce social safety net spending, while Democrats show no preference shifts. In response to AI threat information, those without a degree prefer to regulate AI instead of redistributing wealth to protect against the job threat. These findings present an opportunity for Democrats – and potentially other left-parties globally – to regain their share of the electorate who do not hold a college degree by promising AI regulation. In Chapter 2, I create a cross-national empirical model to test Dani Rodrik's theory that the size and scope of a country's welfare state interacts with the ethnocultural makeup of a society to explain the disparate success of left and right-wing populists as described in 2021's *Why Does Globalization Fuel Populism?* Here, I offer preliminary quantitative evidence to support Rodrik's assertion that left-populists are more successful in homogeneous societies where the ire of the public is more easily directed toward global financial elites seeking to contract the welfare state, while right-populists are more successful when taking advantage of large levels of immigration to campaign on promises of keeping the welfare state exclusive to natives. Finally, Chapter 3 helps us understand how the two dimensions of welfare state breadth – universality and generosity – shape voter preferences toward the regulation of automation technologies that threaten jobs. Welfare states vary widely. The US's welfare state is largely means-tested to ensure benefits only go to those with lower incomes, while others, such as in Scandinavia, favor a universalistic approach to give benefits to all or most citizens. Since welfare states are a form of insurance against job loss, workers should be more comfortable with the integration of job-threatening automation technologies into the economy if they know they will be generously compensated by the welfare state in the event of unemployment. By running a cross-national empirical analysis, I present limited evidence that welfare state universality – not generosity – reduces support for the regulation of automation technology. In sum, this dissertation offers policymakers and analysts a clear picture of the politics of the welfare state in the context of the rapid development of automation technology like artificial intelligence.

Contents

Abstract **ii**

1 Chapter 1: Buying Support for Automation? A Test of the Embedded Liberalism Thesis: Evidence from Europe (2014-2024) **1**

2 Chapter 2: Immigration, the Welfare State, and Explaining Support for Left or Right Populism in Response to Economic Shocks **28**

3 Chapter 3: Compensate the AI Losers?" A Video Experiment on AI-Framing and Policy Preferences **50**

4 Appendix **89**

Buying Support for Automation? A Test of the Embedded Liberalism Thesis: Evidence from Europe (2014-2024)

Julian Haddad

Abstract

I examine whether welfare state generosity can reduce support for regulation of automation technology by extending the Embedded Liberalism thesis beyond trade to also include technological change. While existing literature suggests that social safety nets reduce opposition to globalization, little is known about their effect on public attitudes toward regulating automation technologies which can cause similar labor market disruptions to globalized trade. I conduct a cross-national analysis of automation regulation and automation concern among 24 European countries, comparing countries with different levels of welfare state generosity. Using country-level data from the Comparative Welfare Entitlements Dataset (CWED) and individual-level survey data from Eurobarometer, I assess whether welfare state generosity and universality decrease support for regulating automation technology while controlling for relevant individual level characteristics such as occupation-level risk from AI. I find that AI risk, age, and education are associated with a lower appraisal of automation technologies and the desire to manage them carefully. I do not find strong evidence to suggest that welfare state generosity, however, makes workers desire less regulation of automation. On the contrary, I show that a country's welfare state generosity predicts less concern about the careful management of automation technology among residents and has no relationship with support for AI regulation. Future research should explore welfare state universality as a potential mitigator of automation concern, as I find limited evidence suggesting residents of more universalistic welfare states are less supportive of AI and robotics regulations.

Introduction

Can countries buy support for automation through welfare spending? Automation advancements can disrupt labor markets, and scholars debate whether political backlashes against labor market displacements can be prevented through generous social safety nets (Acemoglu and Johnson, 2023; Acemoglu and Restrepo, 2020). In the context of trade, the Embedded Liberalism literature suggests that countries can "buy" support for trade globalization among displaced workers with welfare state spending via unemployment insurance (Hays et al., 2005; Ritchie and You, 2021). While some scholarship supports this conclusion, other work calls it into question, largely in the context of the political backlash against globalization in the form of anti-trade populist victories (Broz et al., 2021; Colantone and Stanig, 2018). As governments grapple with the potential of massive labor market displacement from advances in automation technology, namely robots and Artificial Intelligence (AI), evaluating whether or not the Embedded Liberalism bargain applies in the context automation-driven displacement is of clear political import.

To answer this question, I conducted a cross-national analysis of citizens' attitudes toward the regulation of automation technologies across a variety of welfare state contexts in Europe. In this study, I distinguish between two key metrics of a welfare state's breadth: generosity and universality. Welfare state generosity refers to the monetary value of social welfare benefits and is measured with data from Scruggs et al. (2017)'s Comparative Welfare Entitlements Dataset (CWED). Welfare state universality refers to the extent to which welfare benefits are available to all citizens, rather than a system where only those below a certain income threshold qualify (Coppedge et al., 2022). Using individual level survey data from the 2014, 2017, and 2024 Eurobarometer special reports on attitudes toward automation and controlling for occupation-level AI risk, I make a number of contributions to the political economy of automation (European Commission And European Parliament, Brussels, 2018; European Commission, 2022; European Commission, 2025; Loaiza and Rigobon, 2024).

Most importantly, I find that citizens in states with greater welfare state generosity do not exhibit greater openness to automation, though I find limited evidence that welfare state universality, may decrease voters' interest in AI regulation. These findings suggest applying the Embedded Liberalism bargain to the labor market displacements portended by 21st century automation may be ineffective at preventing political backlash (Hays et al., 2005; Rodrik, 2021). I state this cautiously, however, and future research should apply this model outside of Europe to see if this conclusion holds. One potential explanation for these findings may be that political attitudes toward automation are more often driven by the fear of price shocks, not job loss, as Magistro et al. (2025) suggest.

Secondly, I contribute to the political economy of automation literature by further illuminate the microfoundations of the politics of AI regulation (Busemeyer and Sahm, 2022; Buser et al., 2020; Green, 2023; Magistro et al., 2025). Those on the political left and with higher levels of education are more supportive of automation regulation and are more likely to have negative attitudes toward AI. In contrast with research from before the rapid proliferation of AI in the form of Large language Models (LLMs), age and education are also associated with support for automation regulation (Golin and Rauh, 2022). This particular finding may reflect the extent to which educated workers are newly threatened by advancements in artificial intelligence, and dovetails with Heinrich and Witko (2025)'s finding that those with greater knowledge about AI are often those who want to regulate it.

Understanding whether generous welfare states might make citizens more comfortable with automation is an important question because the political and economic implications of a country's openness to automation are manyfold. Automation can have profound impacts on a country's long-term economic growth (Prettner, 2019), income inequality (Han, 2022) the risk of its citizens losing their jobs (Frey and Osborne, 2017), and voters' support for populist politicians (Golin and Rauh, 2022; Rodrik, 2021). If the Embedded Liberalism thesis applies to automation technologies in the 21st century, we should expect that citizens of generous welfare states will be more comfortable with automation technologies because

their country's social safety net will protect them in the event of automation-induced job loss. If it does not, we may see a political backlash against the adoption of these technologies reminiscent of the anti-trade-globalization movement of the 2010's and 2020's.

Political Backlashes Against Job Loss and the Demand for Social Safety Nets

Existing research typically views policy demands for insurance in the form of social welfare spending as downstream of individual circumstances (Ohtake and Tomioka, 2004). For example, a person facing a higher chance of job loss should be more supportive of policies that can soften its blow. A body research suggests that individuals specifically prefer more generous social safety nets in the face of increased automation risk. Busemeyer and Sahm (2022) find that individuals in occupations with higher levels of automation risk are more supportive of income redistribution, while Thewissen and Rueda (2019) show that individuals with higher incomes are even more supportive of redistribution in the face of automation threats. The latter finding supports the notion that individuals view social safety nets as a form of insurance against financial loss, wherein those with high incomes and more to lose are even more invested in that insurance.

Alternatively, one can view the welfare state that an individual voter has access to key variable in determining their policy preferences. The existing safety net they have access to, in other words, helps determine the economic policies they are open to. The Embedded Liberalism thesis posits that the existence of a generous welfare state can ameliorate workers to the labor market disruptions of trade by offering them adequate insurance via social spending like unemployment benefits. Globalization in the form of decreased barriers to trade is often disruptive to workers in communities facing newfound competition, like an American auto worker facing potential factory closure due to offshoring of manufacturing or increased competition from Japanese car imports. In terms of pure self-interest, this American worker

should oppose free trade agreements regardless of the promised macroeconomic benefits. Famously, Ruggie (1982) contended that the postwar expansion of the welfare state across western democracies was a "bargain" with workers who would be displaced by the trade liberalization. A number of scholars find empirical evidence for this trend (Cusack et al., 2008) Hays et al. (2005), for example, find that incremental increases in the generosity of unemployment insurance makes workers more open to free trade.

The state of Embedded Liberalism in the 21st century is debated among scholars. In the 2010's, a string of populist victories across the west by candidates who opposed free trade motivated a literature that calls into question the extent to which Embedded Liberalism policy bargain is feasible in the 21st century (Mansfield and Rudra, 2021; Rodrik, 2021). This literature often highlights the ways in which digitally-enabled capital flight made globalization in the 21st century more politically destabilizing than the postwar lowering of trade barriers (Hopkin, 2020). One interpretation of these victories is that the Embedded Liberalism bargain broke down, while another is simply that it welfare benefits need to be better targeted toward the affected workers. Broz et al. (2021) and Colantone and Stanig (2018) find evidence that the communities most dramatically affected by trade liberalization were those where new vote share for anti-trade populists was strongest. A potential implication of these findings is that state investment in the affected communities, rather than simply offering unemployment benefits, may have prevented the political backlash against globalization in the 21st century. Ritchie and You (2021) offers empirical evidence that increased access to redistributive benefits reduced the propensity to vote for Donald Trump in 2016, adding credence to the notion that Embedded Liberalism may still function in the 21st century in the context of trade.

The comparison between automation and trade is often made because both represent capitalistic creative destruction, that can disrupt labor markets while simultaneously providing macroeconomic benefits through increased efficiency and productivity (Magistro et al., 2025). Some evidence suggests that voters' policy preference shifts in the face of automation shocks

are less drastic than trade shocks of the same magnitude (Zhang, 2022). Consider the following results from Di Tella and Rodrik (2020)'s survey experiment where respondents were told that a factory employing 900 people was going to close down due to either: automation, changes in consumer demand, bad management, imports from a poor country, or imports from a rich country. In this case, automation elicits a comparatively muted policy response to trade despite the fact that the economic impact is identical.[1] On the other hand, recent work suggests the advancements of AI is having a greater impact on policy preferences than previous work on automation may have suggested (Green et al., 2025; Heinrich and Witko, 2025; Haslberger et al., 2025)

As workers are increasingly concerned about potential job losses from advances in automation in the form of AI and robotics, it is of clear relevance to policymakers whether the bargain of Embedded Liberalism can be struck. Can an expanded social welfare state "buy" openness to the use of automation technologies like AI? Whether or not a country allows automation technologies to be used freely in the workplace both a policy choice and often a hot-button political issue. Industry regulation of automation technologies, such as limits on the use of AI in media, were a central demand of the recent Writers Guild of America strike, and are reminiscent of protests against factory automation in the 19th century. In the transportation sector, an increasing number of taxi drivers are protesting the threat that driver-less cars pose to their industry, while California Governor, Gavin Newsom, vetoed a bill in 2023 that would have regulated automation technologies in the form of a ban on self-driving trucks in California. Most recently, the Trump Administration issued an executive order asking for AI firms to provide the government early access to new AI models before release.

Busemeyer and Sahm (2022) show that welfare state generosity mediates the degree to which automation threats prime their support for social spending. Citizens of generous welfare states demanded smaller increases in welfare spending as a response to automation

[1]Interestingly, individuals were more responsive to trade shocks from poor countries as opposed to rich ones.

threats, presumably because they already felt they could rely on their existing welfare states. Based on this finding, it is reasonable to surmise that citizens of generous welfare states should be less concerned with regulating automation technologies since they pose less of a threat to one's income. If Embedded Liberalism can be applied in this context, incremental increases in welfare spending should be associated with incremental decreases in worker support for automation regulation in the same way Hays et al. (2005) shows in the context of trade.

Thewissen and Rueda (2019) find that those with higher incomes facing automation threats are even more supportive of redistributive policy responses to automation threats, and the authors argue that this is because they have more to lose. These workers in particular would benefit from a more universalist welfare state, where benefits are given to all regardless of income. Thus, my expansion of the Embedded Liberalism thesis to the context of automation includes welfare state universality as a key metric of welfare state breadth, wherein universality is expected to temper demands for the regulation of automation technologies. Moreover, these policy responses are likely moderated by a number of key personal characteristics such as industry exposure to automation (Thewissen and Rueda, 2019), and gender, with men generally being less concerned about automation (McClure, 2018), education (Golin and Rauh, 2022), age (Green, 2023), and partisanship (Bartels, 2006; Magistro et al., 2025)

Hypothesis 1: Automation exposure increases concern about automation and support for tech regulation.

Hypothesis 2: The impact of automation exposure is smaller when welfare benefits are high.

Data and Methods

Independent Variables: I argue that the extent to which welfare benefits are available is multidimensional, where welfare state generosity is the amount of welfare benefits provided while universality refers to what share of the population has access to welfare benefits. These two dimensions are the main independent variables of interest for this study. I measure welfare state generosity as the post-tax benefit replacement rate for an individual, also known as the Unemployment Insurance Replacement Rate (UIRR). UIRR is a superior measure of welfare state generosity than simply measuring social spending as a percentage of GDP for a number of reasons.

First, social expenditure levels do not differentiate by income level and are strongly correlated with orthogonal trends such as the aging of a country's population (Kvist et al., 2013). A country's social expenditure levels as a percentage of GDP may increase to pay for an aging population, but this does not necessarily mean the welfare state has become more "generous", for example. If Japan begins spending more money on pension payments simply because more people qualify than did in previous decades, their welfare state did not actually become more generous; more people simply qualify for it. UIRR has been successfully operationalized as a measure of welfare state generosity in other work in comparative political economy, such as Scruggs and Hayes (2017). Hays et al. (2005) uses the very similar metric of government spending per unemployed worker in their important empirical test of the Embedded Liberalism thesis with regard to trade. Data on UIRR will come from the Comparative Welfare Entitlements Dataset (CWED), which provides data on unemployment insurance and eligibility for a number of pension and income replacement programs in 33 mostly wealthy countries (Scruggs et al., 2017).

The second independent variable of interest is welfare state universality, which will be measured using V-Dem data. On this dimension, welfare benefits are considered to exist along a spectrum of *means-tested* to *universal*. Means-tested benefits are ones that an individual or family must qualify for, typically because their income is below a certain

threshold. An example of means-tested benefits are the United Kingdom's Universal Credit, which helps residents with living costs. In order to qualify, an individual must have a low income and have less than £16,000 in savings (GOV.UK). A purely universal benefit, on the other hand, is Sweden's *barnbidrag*, a monthly stipend provided to all families with children regardless of their income (Försäkringskassan, 2025) To construct their measure of welfare state universality, V-Dem asked country experts "How many welfare programs are means-tested and how many benefit all (or virtually all) members of the polity?". This variable is an ordinal measure of welfare state universality, with a lower value indicating a more means-tested welfare state and a higher value indicating a more universalistic welfare state (Coppedge et al., 2022).

Dependent Variable: The main dependent variable of interest for this study is support for the regulation of automation technology. Eurobarometer is the main source of these data because of the quality and availability provided, as well as the overlap with more than half of the countries included in the CWED. Eurobarometer's 2014, 2017, and 2024 waves include special reports on respondents' attitudes toward technology. The 2014 and 2017 waves are largely identical, but do not directly ask respondents if they'd like to regulate automation technology. As a kind of proxy, they ask respondents how important it is that robots and AI are carefully managed. Additionally, they ask respondents whether they generally feel positively or negatively toward robots and AI.

The 2024 Eurobarometer wave asks many of the same questions, but also asks four questions about respondents' support toward regulating AI specifically (such as banning AI from making autonomous decisions or requiring workers to be involved in the decision-making process when companies adopt AI). Additionally, the timing of Eurobarometer 2024 is of particular benefit because it is after the widespread proliferation of LLMs like Chat-GPT and offers an interesting point of comparison to the previous waves. I construct a support for AI regulation index from the aforementioned questions.

Occupation-Level Automation Risk and Controls: Occupation level automation

risk is an important control that I measure both subjectively and objectively. Eurobarometers 2014 and 2017 measure *subjective automation risk* by asking how likely it is that a respondent's job could be done by a robot or AI in the future. Eurobarometer 2024, interestingly, does not ask this question. Objective automation risk for all waves is measured by Loaiza and Rigobon (2024)'s EPOCH dataset which includes a continuous AI-risk score ranging from 0 to 1, with 1 being the most at risk of being substituted by AI. I measure risk from being replaced by robots via Webb (2020)'s *percent robot* variable which indicates the percentage of a job that can be replaced by a robot.

Both datasets are automatically matched to SOC-2018 occupation codes, which I then manually mapped to Eurobarometer's 18 occupation categories. It should be noted that these categories do not always perfectly match SOC-2018 categories, and that future research would benefit from individual level data that uses SOC or similar classifications. The full crosswalk I created can be found in the Online Appendix. Eurobarometer provides a respondent's education level, however, this is done by listing the age the respondent left school. Though not a perfect metric of educational attainment, it helps standardize education levels across countries. Data for the GDP per capita of the respondent's home country come from the World Bank.

To estimate the effect of welfare state generosity on citizen openness to automation, I use a Hierarchical Linear Model (HLM). Given that individuals are nested within countries, HLM allows for the assessment of both individual-level predictors (such as occupational risk, education level, and gender) and country-level predictors (welfare generosity, universality, GDP per capita, etc.). This approach ensures that I accurately model variation at both levels while allowing for cross-level interactions (Steenbergen and Jones, 2002). One concern with HLMs in comparative work is that a small number of countries can weaken causal inference (Steenbergen and Jones, 2002). With 24 countries in my analysis, this analysis is on the border of clearing the concerns raised by Steenbergen and Jones (2002) in his landmark work. For this reason, I also show the results of OLS models.

Hierarchical Linear Model (HLM) Specification

The HLM consists of two levels:

Level 1: Within-Country Variation

At the individual level, automation concern is modeled as a function of key demographic and economic predictors:

$$Y_{itj} = \beta_{0j} + \beta_1 \text{Ideology}_{itj} + \beta_2 \text{Education}_{itj} + \beta_3 \text{IndustryExposure}_{itj} + \beta_4 \text{Gender}_{itj} + \beta_5 \text{Age}_{itj} + \epsilon_{itj} \quad (1)$$

where:

- Y_{itj} = Support for Automation Regulation for individual i in country j at time t.
- β_{0j} = Country-specific intercept, representing the baseline level of support for regulating automation within country j.
- ϵ_{itj} = Individual-level error term.

Level 2: Between-Country Variation

At the country level, I model how welfare state generosity and economic conditions shape automation concern:

$$\beta_{0j} = \gamma_0 + \gamma_1 \text{WelfareGenerosity}_{jt} + \gamma_2 \text{Per Capita GDP}_{jt} + \gamma_3 \text{WelfareUniversality}_{jt} + u_j \quad (2)$$

where:

- γ_0 = Overall intercept across all countries.
- γ_1 = Effect of welfare generosity on support for regulating automation at the country level.

- γ_2 = Effect of GDP per capita on support for regulating automation.
- γ_3 = Effect of welfare state universality on support for regulating automation.
- u_j = Random country-level error term, capturing country-specific variation.

For Eurobarometer waves 2014 and 2017, I lag welfare state generosity and welfare state universality by 1 year, which helps identify changes in automation concern that occur after a change in welfare state policy. For analysis of the 2024 wave, I use the UIRR value from 2022 because it is the latest year for which the CWED have available data.

Results

I present the results of four model specifications using both OLS and HLM, which leverage different metrics of support for the regulation of automation. **Tables 1 and 2** include data from Eurobarometer waves 2014 and 2017, but not 2024. This is because the 2014 and 2017 waves asked respondents for their personal appraisal of whether or not their job could be replaced by a robot, but the 2024 wave did not, which does not allow me to control for subjective risk evenly across the sample. **Table 3** includes data from the 2024 Eurobarometer only because only the 2024 wave *explicitly* asked respondents how much they support regulating AI via a number of AI regulation questions. **Table 4** shows a pooled analysis of all waves, wherein subjective automation risk is removed because of its missingness in 2024. I supplement this specification with a measure of AI risk from Loaiza and Rigobon (2024) and a measure of robot risk from Webb (2020)

Let us begin with Tables 1 and 2. The dependent variable used in Table 1 is the extent to which respondents believe that "Robots and artificial intelligence are technologies that require careful management." It should be noted that this question does not explicitly get at whether this management ought to be through government regulation, though it is a reasonable assumption a respondent may make.

Table 1: Welfare State Generosity and Support for the Careful Management of AI (2014 2017)

	OLS	HLM
(Intercept)	3.119***	2.751***
	(0.105)	(0.134)
Subjective Robot Risk	−0.009	−0.004
	(0.008)	(0.006)
UIRR (Welfare Generosity)	0.291	0.794***
	(0.178)	(0.180)
Welfare Universality (V-Dem)	−0.042	0.038
	(0.049)	(0.050)
Left-Right Orientation	−0.008**	−0.007***
	(0.004)	(0.002)
AI Risk	0.211**	0.236***
	(0.085)	(0.073)
Robot Risk	−0.001**	−0.001***
	(0.000)	(0.000)
Age	0.003***	0.002***
	(0.001)	(0.000)
Education (age left school)	0.001	−0.001
	(0.002)	(0.001)
Female	0.028*	0.028***
	(0.017)	(0.010)
Wave: 2017	0.030	0.024**
	(0.032)	(0.011)
GDP per capita	0.003**	0.003
	(0.001)	(0.002)
SD (Intercept iso3)		0.158
SD (Observations)		0.678
N (individuals)	18 907	18 907
ICC		0.052
Random intercept variance	—	0.025
N countries (Level 2)	24	24
Estimation	OLS	REML

* p <0.1, ** p <0.05, *** p <0.01
DV: Careful Management of AI (1–4).

Table 1 does not offer support for either Hypothesis 1 or 2. Rather, it shows that welfare state generosity is in fact associated with greater support for the careful management of robots and AI while welfare state universality has no relationship with it. One potential explanation for this is that welfare state generosity is correlated with other macro-level characteristics of a country that may predict support for regulation. For example, it may be associated with the extent to which voters are educated about emerging technologies and concerned about their potential impacts. Recall that this table measures respondents' beliefs in 2014 and 2017, a time when AI was not understood to be a commonly used technology and robots were largely only integrated into the economy in factory automation settings.

In congruence with previous research, women and those on the political left support greater management of AI and robots. AI risk as measured by Loaiza and Rigobon (2024) is associated with greater support for careful management of AI and robots, but robot risk is not. Lastly, the 2017 wave saw an increase in support for carefully managing robots and AI.

Table 2 includes the same set of independent variables but, as its dependent variable, simply measures how positively a person views robots. Welfare state generosity as measured by UIRR is not associated with a more positive appraisal of robots, but welfare state universality is. This finding offers some support for Hypothesis 2. All other control variables move in expected directions when compared to previous research (e.g., women, older individuals, and those in occupations with higher robot/AI risk view robots more negatively). Counterintuitively, those who believe that their own job could be done by a robot (labeled Subjective Robot Risk) have more positive appraisals of robots.

There are a number of limitations to these first two models' specifications, though they offer useful insights. The first is that respondents' occupations is measured by Eurobarometer via 18 categories for which do not exist a crosswalk to widely used occupation classifications such as SOC-2018, which most metrics of automation risk map to. Because of this, I manually mapped the 18 Eurobarometer occupation categories to this dataset, however, some

Table 2: General Robot Appraisal (2014 2017)

	OLS	HLM
(Intercept)	2.715***	2.766***
	(0.124)	(0.136)
Subjective Robot Risk	0.126***	0.128***
	(0.016)	(0.006)
UIRR (Welfare Generosity)	0.065	−0.083
	(0.211)	(0.181)
Welfare Universality (V-Dem)	0.058	0.124**
	(0.065)	(0.051)
Left-Right Orientation	0.010*	0.004*
	(0.006)	(0.002)
AI Risk	−0.197*	−0.168**
	(0.119)	(0.077)
Robot Risk	−0.003***	−0.003***
	(0.000)	(0.000)
Age	−0.002***	−0.003***
	(0.001)	(0.000)
Education (age left school)	0.006***	0.002***
	(0.002)	(0.001)
Female	−0.131***	−0.134***
	(0.015)	(0.010)
Wave: 2017	−0.060***	−0.064***
	(0.018)	(0.011)
GDP per capita	0.003	0.004**
	(0.002)	(0.002)
SD (Intercept iso3)		0.156
SD (Observations)		0.709
N (individuals)	18 339	18 339
ICC		0.046
Random intercept variance	—	0.024
N countries (Level 2)	24	24
Estimation	OLS	REML

* p <0.1, ** p <0.05, *** p <0.01
DV: Robot Appraisal (1=very neg., 4=very pos.). Reference wave: 2014. OLS SEs clustered by country; random intercepts by country.

categories are quite ambiguous and, thus, AI risk and robot risk as independent variables may not fully capture what they intend to.

Finally, neither question used as dependent variables explicitly asks about the regulation of automation technology by the state. Table 1 gets closer to this question by asking about support for the careful management of robots and AI, but it is worth noting that a respondent might believe these ought to be carefully managed by the companies that run them or by individuals in their personal lives. The 2024 Eurobarometer wave asks four questions about AI regulation, specifically, and thus Table 3 shows a model specified solely on Eurobarometer's 2024 data with an index of respondents' answers to those regulation questions as the dependent variable.

Table 3 shows that welfare state generosity and universality are not correlated with decreases in support for AI regulation in HLM models, however, the OLS specification does show welfare state universality to be associated with less support for AI regulation. Given that HLM models are better suited for nested data analysis, I do not consider this to offer strong support for either hypothesis. Robot risk, interestingly, is associated with *less* support for AI regulation, but not AI risk, while education, age, and left leaning politics are associated with stronger support for AI regulation in 2024. Finally, I construct a pooled analysis of all waves in Table 4. Because the only consistently asked question across all three waves is the extent to which respondents believe AI and robots should be carefully managed, I use this as my dependent variable. Subjective robot risk is dropped from this analysis because it does not appear in the 2024 Eurobarometer wave.

Counter to expectations, welfare state generosity is associated with greater support for carefully managing robots and AI. In keeping with previous literature, those on the political left and those with greater AI risk are more supportive of carefully managing AI and robots (Green, 2023). When compared with 2014, those in 2017 were more concerned with carefully managing the technology while those in 2024 were *less* concerned with it.

Table 3: Welfare Generosity and Support for AI Regulation (2024)

	OLS	HLM
(Intercept)	3.066***	3.057***
	(0.151)	(0.152)
UIRR (Welfare Generosity)	0.259	0.245
	(0.167)	(0.161)
Welfare Universality (V-Dem)	−0.075*	−0.078
	(0.040)	(0.052)
Left-Right Orientation	−0.014***	−0.013***
	(0.005)	(0.003)
AI Risk	0.073	0.125
	(0.123)	(0.094)
Robot Risk	−0.003***	−0.003***
	(0.001)	(0.000)
Age	0.002**	0.001***
	(0.001)	(0.001)
Education (age left school)	0.004	0.003***
	(0.003)	(0.001)
Female	−0.005	−0.004
	(0.017)	(0.013)
GDP per capita	0.004***	0.005***
	(0.001)	(0.001)
SD (Intercept iso3)		0.109
SD (Observations)		0.659
N (individuals)	10 697	10 697
ICC		0.027
Random intercept variance	—	0.012
N countries (Level 2)	23	23
Estimation	OLS	REML

* p <0.1, ** p <0.05, *** p <0.01
DV: AI Regulation Index (1–4). 2024 only. OLS SEs clustered by country; random intercepts by country.

Table 4: Welfare Generosity, Universality, and Careful Management of AI. All Waves Pooled (2014, 2017, 2024)

	OLS	HLM
(Intercept)	3.088***	3.121***
	(0.103)	(0.098)
UIRR (Welfare Generosity)	0.365**	0.307**
	(0.163)	(0.137)
Welfare Universality (V-Dem)	−0.053	−0.008
	(0.039)	(0.023)
Left-Right Orientation	−0.009**	−0.008***
	(0.003)	(0.002)
AI Risk	0.095	0.120**
	(0.093)	(0.059)
Robot Risk	−0.002***	−0.001***
	(0.001)	(0.000)
Age	0.003***	0.002***
	(0.001)	(0.000)
Education (age left school)	0.003***	0.001
	(0.001)	(0.001)
Female	0.004	0.004
	(0.014)	(0.008)
Wave: 2017	0.024	0.024**
	(0.031)	(0.010)
Wave: 2024	−0.080*	−0.070***
	(0.049)	(0.011)
GDP per capita	0.004***	0.004***
	(0.001)	(0.001)
SD (Intercept iso3)		0.119
SD (Observations)		0.699
N (individuals)	29 959	29 959
Random intercept variance	—	0.014
N countries (Level 2)	24	24
Estimation	OLS	REML

DV: Careful Management of AI (1–4). All waves pooled.

Discussion

The Embedded Liberalism hypothesis posited that the political backlash against labor market disruptions from global trade could be prevented through the provision of generous social safety nets for affected workers (Hays et al., 2005). However, the validity of these assumptions has been called into question by the right-wing populist backlash against the global trade order that has taken root in some of the countries where Embedded Liberalism was championed by policymakers (Mansfield and Rudra, 2021). On the other hand, some studies have found support for populists as recently as 2016 was mitigated by increased redistribution (Ritchie and You, 2021). As electorates are grappling with the extent to which they wish to regulate emerging automation technologies like AI, the relevance evaluating the Embedded Liberalism hypothesis in the 21st century is of clear utility. It is important to consider the extent to which institutional factors like the width and depth of a country's welfare state, so to speak, may influence how voters wish for the government to treat technologies that could threaten their livelihoods. If the Embedded Liberalism thesis can be mapped onto voters' worries about automation, one might expect that a deregulated AI industry, for example, may be most championed by citizens of generous and universalist welfare states like Sweden because they have less to lose in the case of a job loss.

I find evidence to cautiously suggest that the Embedded Liberalism hypothesis does *not* hold with relation to automation technologies like robots and AI. In contrast to expectations, welfare state generosity across European economies associated with more concern about carefully managing AI and robots over a ten year time-span. Only in one model specification (OLS across the 2024 Eurobarometer wave, but not HLM) does the relationship between the welfare state and automation regulation move in the expected direction. In that case, welfare state universality – the extent to which welfare benefits are universally available vs. means-tested – predicts less support for AI regulation.

One reason welfare state generosity may not be associated with an electorate friendly to automation technologies may be that unemployment insurance doesn't provide the psy-

chological benefits of being employed (Bénabou and Tirole, 2006, 2016; Lerner, 1980). This same conclusion could be drawn with regard to the relationship between social safety nets and support for global trade among those affected, which might explain the conflicting results in the literature as to the efficacy of the Embedded Liberalism bargain in that context. It is also entirely possible that the dimensions of voter behavior in this context are more multidimensional than is gotten at by studies like the present one. For example, Broz et al. (2021) suggests that voters respond in particular to the ways in which economic shocks affect their communities. If this is also the case with automation technologies, then we might expect welfare state generosity only to matter to those in communities disproportionately affected by labor market displacement. Complicating the picture still, new research suggests that voters' concerns about AI and other forms of automation may be activated if and when they see inflation caused by these technologies in their everyday lives (Magistro et al., 2025).

Given that this analysis is limited to 23 countries across Europe, I caution against drawing strong conclusions based on these results. Future research should, ideally, use global data on both welfare state generosity and support for regulating automation technology to confirm whether the Embedded Liberalism bargain holds in the context of automation in the 21st century. If it does, policymakers can prevent a political backlash against AI through welfare state expansion. If not, they may need to steel themselves for another wave of political backlash against the backdrop of a rapidly changing economy.

References

Daron Acemoglu and Simon Johnson. *Power and Progress: our thousand-year struggle over technology and prosperity*. 2023.

Daron Acemoglu and Pascual Restrepo. Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244, June 2020. ISSN 0022-3808. doi: 10.1086/705716. URL https://www.journals.uchicago.edu/doi/10.1086/705716.

Larry M Bartels. What's the Matter with What's the Matter with Kansas? 2006.

J. Lawrence Broz, Jeffry Frieden, and Stephen Weymouth. Populism in Place: The Economic Geography of the Globalization Backlash. *International Organization*, 75(2):464–494, February 2021. ISSN 0020-8183, 1531-5088. doi: 10.1017/S0020818320000314. URL https://www.cambridge.org/core/journals/international-organization/article/populism-in-place-the-economic-geography-of-the-globalization-backlash/98ED873D925E0590CB9A78AEC68BB439.

Marius R. Busemeyer and Alexander H. J. Sahm. Social Investment, Redistribution or Basic Income? Exploring the Association Between Automation Risk and Welfare State Attitudes in Europe. *Journal of Social Policy*, 51(4):751–770, October 2022. ISSN 0047-2794, 1469-7823. doi: 10.1017/S0047279421000519.

Thomas Buser, Gianluca Grimalda, Louis Putterman, and Joël van der Weele. Overconfidence and gender gaps in redistributive preferences: Cross-Country experimental evidence. *Journal of Economic Behavior & Organization*, 178:267–286, October 2020. ISSN 0167-2681. doi: 10.1016/j.jebo.2020.07.005. URL https://www.sciencedirect.com/science/article/pii/S0167268120302274.

Roland Bénabou and Jean Tirole. Belief in a Just World and Redistributive Politics*. *The Quarterly Journal of Economics*, 121(2):699–746, May 2006. ISSN 0033-5533. doi: 10.1162/qjec.2006.121.2.699. URL https://doi.org/10.1162/qjec.2006.121.2.699.

Roland Bénabou and Jean Tirole. Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, 30(3):141–164, August 2016. ISSN 0895-3309. doi: 10.1257/jep.30.3.141. URL https://pubs.aeaweb.org/doi/10.1257/jep.30.3.141.

Italo Colantone and Piero Stanig. Global Competition and Brexit. *American Polit-*

ical Science Review, 112(2):201–218, May 2018. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055417000685. URL https://www.cambridge.org/core/journals/american-political-science-review/article/global-competition-and-brexit/C843990101DB9232B654E77130F88398.

Michael Coppedge, John Gerring, and Carl Henrik. Dataset Archive – V-Dem, 2022. URL https://www.v-dem.net/data/dataset-archive/.

Thomas R. Cusack, Torben Iversen, and Philipp Rehm. Economic Shocks, Inequality, and Popular Support for Redistribution. In Pablo Beramendi and Christopher J. Anderson, editors, *Democracy, Inequality, and Representation in Comparative Perspective*, pages 203–231. Russell Sage Foundation, 2008. ISBN 978-0-87154-088-1. URL http://www.jstor.org/stable/10.7758/9781610440448.11.

Rafael Di Tella and Dani Rodrik. Labour Market Shocks and the Demand for Trade Protection: Evidence from Online Surveys. *The Economic Journal*, 130(628):1008–1030, May 2020. ISSN 0013-0133. doi: 10.1093/ej/ueaa006. URL https://doi.org/10.1093/ej/ueaa006.

European Commission. Eurobarometer 87.1: Two Years Until the 2019 European Elections, Attitudes of Europeans Towards Tobacco and Electronic Cigarettes, Climate Change, Attitudes Towards the Impact of Digitization and Automation on Daily Life, and Coach Services, March 2017, 2022.

Brussels European Commission. Eurobarometer 101.4 (2024), 2025. Published: (ZA8844; Version 1.0.0) [Data set]. GESIS, Cologne.

European Commission And European Parliament, Brussels. Eurobarometer 82.4 (2014)Eurobarometer 82.4 (2014): Parlemeter 2014: Parlemeter 2014: The European Parliament, Autonomous Systems, Gender Equality, and Smoking Habits: The European

Parliament, Autonomous Systems, Gender Equality, and Smoking Habits, 2018. URL https://search.gesis.org/research_data/ZA5933?doi=10.4232/1.13044.

Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280, January 2017. ISSN 0040-1625. doi: 10.1016/j.techfore.2016.08.019. URL https://www.sciencedirect.com/science/article/pii/S0040162516302244.

Försäkringskassan. Child allowance, July 2025. URL https://www.forsakringskassan.se/english/parents/child-allowance.

M. Golin and C. Rauh. The Impact of Fear of Automation. *Janeway Institute Working Papers*, December 2022. URL https://ideas.repec.org//p/cam/camjip/2229.html. Number: 2229.

GOV.UK. Universal Credit. URL https://www.gov.uk/universal-credit/eligibility.

Jane Green, Zack Grant, Geoffrey Evans, and Gaetano Inglese. Linking artificial intelligence job exposure to expectations: Understanding AI losers, winners, and their political preferences. *Research & Politics*, 12(2):20531680251337897, April 2025. ISSN 2053-1680. doi: 10.1177/20531680251337897. URL https://doi.org/10.1177/20531680251337897.

Scott Keeter and Ted Van Hannah Hartig Green, Andrew Daniller. Demographic profiles of Republican and Democratic voters, July 2023. URL https://www.pewresearch.org/politics/2023/07/12/demographic-profiles-of-republican-and-democratic-voters/.

Yang Han. The Impact of Industrial Robots on the Skill-Based Wage Gap. *American Journal of Industrial and Business Management*, 12(4):571–602, April 2022. doi: 10.4236/ajibm.2022.124031. URL https://www.scirp.org/journal/paperinformation?paperid=116642. Number: 4.

Matthias Haslberger, Jane Gingrich, and Jasmine Bhatia. Rage against the machine? Generative AI exposure, subjective risk, and policy preferences. *Journal of European Public Policy*, 0(0):1–28, September 2025. ISSN 1350-1763. doi: 10.1080/13501763.2025.2554903. URL https://doi.org/10.1080/13501763.2025.2554903. _eprint: https://doi.org/10.1080/13501763.2025.2554903.

Jude C. Hays, Sean D. Ehrlich, and Clint Peinhardt. Government Spending and Public Support for Trade in the OECD: An Empirical Test of the Embedded Liberalism Thesis. *International Organization*, 59(2):473–494, April 2005. ISSN 1531-5088, 0020-8183. doi: 10.1017/S0020818305050150.

Tobias Heinrich and Christopher Witko. Self-interest and preferences for the regulation of artificial intelligence. *Journal of Information Technology & Politics*, 22(3):306–321, July 2025. ISSN 1933-1681. doi: 10.1080/19331681.2024.2370815. URL https://doi.org/10.1080/19331681.2024.2370815. _eprint: https://doi.org/10.1080/19331681.2024.2370815.

Jonathan Hopkin. Introduction: A Quick History of the Present. In Jonathan Hopkin, editor, *Anti-System Politics: The Crisis of Market Liberalism in Rich Democracies*, page 0. Oxford University Press, June 2020. ISBN 978-0-19-069976-5. doi: 10.1093/oso/9780190699765.003.0001. URL https://doi.org/10.1093/oso/9780190699765.003.0001.

Jon Kvist, Simon Grundt Straubinger, and Anders Freundt. Measurement Validity in Comparative Welfare State Research: The Case of Measuring Welfare State Generosity. *European Journal of Social Security*, 15(4):321–340, December 2013. ISSN 1388-2627. doi: 10.1177/138826271301500402. URL https://doi.org/10.1177/138826271301500402.

Melvin J. Lerner. The Belief in a Just World. In Melvin J. Lerner, editor, *The Belief in a Just World: A Fundamental Delusion*, pages 9–30. Springer US, Boston, MA, 1980. ISBN

978-1-4899-0448-5. doi: 10.1007/978-1-4899-0448-5_2. URL https://doi.org/10.1007/978-1-4899-0448-5_2.

Isabella Loaiza and Roberto Rigobon. The EPOCH of AI: Human-Machine Complementarities at Work, November 2024. URL https://papers.ssrn.com/abstract=5028371.

Beatrice Magistro, Sophie Borwein, R. Michael Alvarez, Bart Bonikowski, and Peter John Loewen. Attitudes toward artificial intelligence (AI) and globalization: Common microfoundations and political implications. *American Journal of Political Science*, 70(1):348–365, 2025. ISSN 1540-5907. doi: 10.1111/ajps.12959. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ajps.12959. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12959.

Edward D. Mansfield and Nita Rudra. Embedded Liberalism in the Digital Era. *International Organization*, 75(2):558–585, 2021. ISSN 0020-8183. URL https://www.jstor.org/stable/27104649.

Paul K. McClure. "You're Fired," Says the Robot: The Rise of Automation in the Workplace, Technophobes, and Fears of Unemployment. *Social Science Computer Review*, 36(2): 139–156, April 2018. ISSN 0894-4393. doi: 10.1177/0894439317698637. URL https://doi.org/10.1177/0894439317698637.

Fumio Ohtake and Jun Tomioka. Who Supports Redistribution? *The Japanese Economic Review*, 55(4):333–354, December 2004. ISSN 1468-5876. doi: 10.1111/j.1468-5876.2004.00318.x. URL https://doi.org/10.1111/j.1468-5876.2004.00318.x.

Klaus Prettner. A NOTE ON THE IMPLICATIONS OF AUTOMATION FOR ECONOMIC GROWTH AND THE LABOR SHARE. *Macroeconomic Dynamics*, 23(3):1294–1301, April 2019. ISSN 1365-1005, 1469-8056. doi: 10.1017/S1365100517000098.

Melinda N. Ritchie and Hye Young You. Trump and Trade: Protectionist Politics and

Redistributive Policy. *The Journal of Politics*, 83(2):800–805, April 2021. doi: 10.1086/710322. URL https://www.journals.uchicago.edu/doi/full/10.1086/710322.

Dani Rodrik. Why Does Globalization Fuel Populism? Economics, Culture, and the Rise of Right-Wing Populism. *Annual Review of Economics*, 13(Volume 13, 2021):133–170, August 2021. ISSN 1941-1383, 1941-1391. doi: 10.1146/annurev-economics-070220-032416. URL https://www.annualreviews.org/content/journals/10.1146/annurev-economics-070220-032416.

John Gerard Ruggie. International Regimes, Transactions, and Change: Embedded Liberalism in the Postwar Economic Order. *International Organization*, 36(2):379–415, 1982. ISSN 0020-8183. URL https://www.jstor.org/stable/2706527.

Lyle Scruggs and Thomas J. Hayes. The Influence of Inequality on Welfare Generosity: Evidence from the US States. *Politics & Society*, 45(1):35–66, March 2017. ISSN 0032-3292. doi: 10.1177/0032329216683165. URL https://doi.org/10.1177/0032329216683165.

Lyle Scruggs, Jahn Detlef, and Kati Kuitto. Comparative Welfare Entitlements Dataset (CWED2), 2017. URL http://cwed2.org/.

Marco R. Steenbergen and Bradford S. Jones. Modeling Multilevel Data Structures. *American Journal of Political Science*, 46(1):218–237, 2002. ISSN 0092-5853. doi: 10.2307/3088424. URL https://www.jstor.org/stable/3088424.

Stefan Thewissen and David Rueda. Automation and the Welfare State: Technological Change as a Determinant of Redistribution Preferences. *Comparative Political Studies*, 52(2):171–208, February 2019. ISSN 0010-4140. doi: 10.1177/0010414017740600. URL https://doi.org/10.1177/0010414017740600.

Michael Webb. The Impact of Artificial Intelligence on the Labor Market. 2020.

Baobao Zhang. No Rage Against the Machines: Threat of Automation Does Not Change Policy Preferences. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 856–866, Oxford United Kingdom, July 2022. ACM. ISBN 978-1-4503-9247-1. doi: 10.1145/3514094.3534179. URL https://dl.acm.org/doi/10.1145/3514094.3534179.

Immigration, the Welfare State, and Explaining Support for Left or Right Populism in Response to Economic Shocks

Julian Haddad

Abstract

The welfare state looms large in populist politics of both a left-wing and right-wing variety, but little empirical work tests the relationship between welfare state contexts and the success of different varieties of populists. While left-populists often argue for the expansion or retention of the existing welfare state in the face of financial crises and calls for austerity, right-populists often engage in welfare chauvinism, wherein immigrants are viewed as putting undue pressure on the welfare state. It is for these reasons that Rodrik (2021) compellingly argues that ethnically diverse societies with generous welfare states should be more amenable to right-wing populist movements in the face of economic shocks. Left-populists, on the other hand, should be more successful in homogeneous societies with weaker welfare states where anti-austerity messages and calls for expanding the welfare state will be more successful when dealing with economic shocks. I test a preliminary, cross-national empirical model using data from the Global Party Survey, World Bank, and Systemic Banking Crises Database II. I offer limited empirical support for the assumptions Rodrik (2021) makes, namely that immigration influxes are associated with greater success of right populists, while more homogeneous societies tend to vote for left populists. Due to limited data availability (n=132), I caution strongly against drawing conclusions from these results and highlight steps for improving causal inference.

Introduction

Electorates are sometimes in "populist moods", but vary in whether they elect left-wing or right-wing populists. What explains why electorates sometimes choose left-populists like Hugo Chavez in Venezuela or Alexis Tsipras in Greece, and other times elect right-populists like Donald Trump in the United States, Giorgia Meloni in Italy, or Narendra Modi in India? This paper begins to answer this question by testing a cross-national empirical model to show the correlates of left or right populist victories across 32 unique election based a modified version of Rodrik (2021)'s theory of left and right populist success. The analysis presented in this study is preliminary, and I outline ways future research in this area can be more robust and offer stronger causal inference Results and Discussion sections.

Populism is an anti-system style of politics characterized by arguments that the existing economic and political order has failed large swathes of the population (Broz et al., 2021; Mudde and Kaltwasser, 2013). While mainstream left and right parties typically differ strongly on economic issues; namely, the size and generosity of the welfare state, this is not the case for populist movements. Generally speaking, populist parties of both the left and the right are often protective of the welfare state regardless of where they stand on social issues (Hopkin, 2020). Broadly speaking, left-wing populists often propose expansions of the welfare state, protection of the welfare state from austerity measures, expropriation of key industries, and combative relationships with large financial institutions. Right-wing populists, on the other hand, typically offer anti-immigration, nationalistic policy agendas; as populists tend to be anti-globalization, there are examples of both left and right-wing populists around the world who have argued against free trade (Hopkin, 2020; Di Tella and Rodrik, 2020; Rodrik, 2021)

Scholars have sought to understand what explains the recent (but not the first) wave global success of *right-wing* populist parties since the 2010's, with a key motivation for this level of interest being the observed tendency for populist leaders to disrespect norms of liberal democracy (to various degrees, and not universally) (Galston, 2018). Currently,

the literature lacks a comprehensive, cross national, empirical model that predicts whether left or right-wing populist parties will be successful. This paper constructs that model by empirically testing a modified version of the framework posited by Rodrik (2021) to include insights from Hopkin (2020), building upon empirical work that largely predicts the success of right-wing populists.

1 The Welfare State and The Choice Between Left & Right Populism

The welfare state plays a large and somewhat complicated role in producing support for both left and right populists, as, in contrast with traditional left-right divides, both left and right populists are typically protective of the welfare state.[1] In a break with economically conservative mainstream parties, many right-populists campaign on arguments that are protective of the welfare state and are characterized by "welfare chauvinism" stating that they will ensure only 'deserving' or 'true' members of the nation actually receive welfare benefits, often to the exclusion of immigrants (Fenger, 2018; Hopkin, 2020). Consider the UK Leave campaign's emphasis on reducing immigration and leaving the EU with the stated goal of funding the NHS (Norris and Inglehart, 2019). Thus, right-wing populists' opposition to immigration is often to protect the welfare state, as natives may view immigrants as putting pressure on the welfare state, and this effect appears to be magnified when the welfare state is larger (Hatton, 2016; Cavaille, 2017).

(Rodrik, 2021) offers a theoretical framework for sorting out exactly how the welfare state plays into this story by arguing that welfare state size interacts with different kinds of globalization shocks to produce support for populists. At the most basic level, he argues that immigration shocks tend to produce support for right-wing populists, especially when

[1] It is worth noting that is not universal, but (Hopkin, 2020) shows that 'anti-system' parties, which are mostly populist parties are indeed typically protective of the welfare state.

there exist strong ethnic and cultural cleavages, while trade and financial shocks produce support for left-wing populists. Thus, high levels of immigration coupled with generous welfare spending should create electoral conditions more advantageous to right-wing populists, while lower levels of welfare spending coupled with trade or financial shocks will produce support for left-populists.

By expanding his theory in keeping with what other scholars have found regarding the welfare state as an institution that impacts political behavior, a couple of key deductions can be made. First, the impact of immigration shocks on right populism should be stronger when welfare spending is high because of concerns for immigrants' impact on the welfare state. Second, the impact of trade/financial shocks on left populism should be weaker when welfare spending is high because voters are better compensated for losses, thus minimizing voter defections (Hays et al., 2005).

While a great deal of research helps us understand when right-populists will be successful cross-nationally (Broz et al., 2021; Colantone and Stanig, 2018; Norris and Inglehart, 2019), little is known about conditions that produce support for left-populists, and there exists no comprehensive empirical model that combines the various theoretical insights in this literature in order to predict the disparate success of left and right wing populists globally. This is an important open question because left-populists produce meaningfully different outcomes for their constituents.

What is so Different About Left Populists?

Populism is considered by most scholars to be a "thin-centered ideology", meaning that it is a style of politics as opposed to a clear set of policy prescriptions (Mudde and Kaltwasser, 2013). However, some patterns emerge that can help distinguish between populism's right and left-wing variants. Right-populists typically offer nationalistic policy agendas aimed at reducing immigration and limiting free trade, while left populists tend to offer expansions of the welfare state, anti-austerity measures, regulations of the financial sector, and sometimes

even industry nationalization; as well as often sharing right-populists' opposition to free trade (Rodrik, 2021).

Though left and right populists often respond to the same grievances (job losses from international trade, national or international financial shocks, a feeling that elites are exploiting the common folk) there are clear differences between the tangible effects of left and right populist governments in both their immigration and social policy. For example, the left populist parties Podemos in Spain and Syriza in Greece railed against austerity, with little of the focus on anti-immigration politics or social conservatism that we saw from right populists like Donald Trump in the US, Nigel Farage in the UK, Jair Bolsonaro in Brazil, or Viktor Orbán in Hungary (Hopkin, 2020). In the wave of populist successes beginning in the 2010's, those of the right-wing variety have been far more successful, and have concerned many analysts due to their propensity for undermining democratic institutions (Galston, 2018). However, I argue that it is a mistake to assume either right-wing populism is the only electorally viable form of populism or 2) that it is the only form of populism that poses a threat to democracy.

Though left populists have not been nearly as successful as right populists globally since the 2010s, left populists have had electoral success and undermined democracy[2]. Hugo Chavez of Venezuela was a quintessential left-populist who came to power in 1999 among a wave of Latin American left-populist election victories in the 1990's and 2000's, sometimes called the "pink tide" (Roberts, 2007) The welfare state typically features heavily in left-populist agendas. Consider Chavez, whose policies differed significantly from what we have come to see from right-populists like Donald Trump or Viktor Orbán in that they were not anti-immigrant in nature. Chavez was a socialist who widely expanded the welfare state (Howard-Hassmann, 2015), opposed an international system that he argued benefited financial elites at the expense of the common people, and nationalized a number of large corporations in key industries such as oil, agriculture, finance, gold, and steel (Reuters, 2012;

[2]In fact, Socrates's famous description of the tyrant in Plato's *Republic* is a left-populist who promises land and wealth redistrubution and ultimately becomes an authoritarian.

Roberts, 2007). He also transformed Venezuela from a democracy into an authoritarian regime, cultivated a cult of personality, and severely undermined human rights (Kornblith, 2013). What was definitively *not* a part of Chavez's tenure was anti-immigrant policy or sentiments in the way that it would be for right-populists.

Of course, other left populists are not anti-democratic. In 2015, Greece was governed by a left-populist Syriza party that maintained a contentious relationship with the European Union in the midst of the Eurozone economic crisis. The Syriza government, buoyed by support from a Greek population struggling with the effects of a debt crisis, was able to delay austerity measures in an effort to protect the welfare state (Aslanidis and Rovira Kaltwasser, 2016). Left-populists are currently a meaningful political force in the United States, as illustrated by the most recent victory of Zohran Mamdani in New York City. However, no research to my knowledge specifically tests what conditions specifically produce support for *left* populist parties globally. To remedy this issue, I create the first full, cross-national predictive model of what explains the disparate success of both left-wing and right-wing populists.

2 Theory

Rodrik (2021) argues that the success of populist politicians can be largely explained by globalization shocks. Specifically, he finds evidence that left-populists become more popular when globalization shocks take the form of financial crises wherein the corrupt elite that becomes the focus of the public's ire is typically a faceless, often international financial elite threatening to gut the existing welfare state with austerity measures. It is this description that fits the left populist wave Latin America saw in the 1990's and 2000's. Right-wing populists, he argues, become popular when globalization entails large immigration shocks, particularly ones that exacerbate ethnonational cleavages which, among other things, may threaten the existing welfare state (i.e., Europe in the face of the Syrian Refugee Crisis in

the 2010's). This phenomenon is often called "Welfare Chauvinism" (Fenger, 2018; Hopkin, 2020). Rodrik, I believe, correctly centers welfare state contexts in explaining differential support for left and right populists.

Rodrik's theory is particularly interesting with regard to what he calls mixed cases that do not fit neatly into the two categorizations offered above, such as the United States in the early 21st century. In the years preceding the 2016 election, the US saw large immigration flows from Latin America (Hanson et al., 2017). Coupled with the US's a meaningful welfare state and ethnic diversity; Rodrik's theory would predict the success of a right populist like Donald Trump (Rodrik, 2021). However, at the same time, the US faced the 2008 Global Financial Crisis, automation shocks to manufacturing (Di Tella and Rodrik, 2020; Acemoglu and Restrepo, 2020), a trade shock from China (Autor et al., 2016), all of which produced support for left-populists like Bernie Sanders. If the US had the predictors of both left and right populism, what explains why a right-populist like Trump was more successful than his left-populist counterparts? Creating a comprehensive model of what predicts left vs. right populist support, as I do, helps to answer this question.

First, it is necessary to identify which portions of Rodrik (2021)'s largely qualitative theory survive empirical scrutiny already conducted in the literature. In a finding that may contradict Rodrik (2021)'s argument that financial globalization shocks produce support for left populists, Funke et al. (2016) find that far-right parties – not left populists, as Rodrik predicts – increased their vote share after financial crises from 1870-2014 across advanced democracies. However their analysis includes *all* banking crises, not just those caused by globalization shocks, so it is unclear whether their findings truly contradict Rodrik (2021)'s core thesis which privileges globalization shocks. Additionally, Funke et al. (2016) does not control for different levels of welfare state generosity, immigration levels, or ethnic fractionalization which, according to Rodrik (2021)'s framework, should impact left and right populist vote shares. Thus, this study works under the assumption that Funke et al. (2016)'s findings do not conclusively disprove Rodrik's claims regarding the welfare state and support

for left populists.

Hypothesis 1: *Globalization shocks that take the form of financial crises will increase support for left-populist parties.*

Rodrik (2021) compellingly argues that the comparatively less intense populist wave in Europe as compared to the United States is due to the size and generosity of European welfare states. Welfare states are credible, long term signals of a government's willingness to compensate those losing out from trade shocks and open markets that can not easily be reneged upon, which some research has shown to mitigate political backlashes against trade (Hays et al., 2005; Ruggie, 1982). While many US trade deals promise compensation for those now competing against imports, that promised compensation often falls short or is nonexistent. Coupled with the US's welfare state being weaker than most European economies, support for populists should generally be greater in the US. Rodrik (2021) theorizes and offers limited empirical evidence that the comparative lack of cultural and religious cleavages in Latin America when compared to the US favored left-wing populist movements, as the out-group hostility was directed at financial elites as opposed to immigrants.[3]

Hypothesis 1a: *This effect will be moderated by the size and generosity of the welfare state.*

Hypothesis 1c: *Smaller social/cultural cleavages and lower levels of immigration will lend support to left populists.*

A number of studies corroborate Rodrik (2021) claim about immigration shocks producing support for right-wing populists(Alesina et al., 2001; Colantone and Stanig, 2018), and this may be because natives often view immigrants as putting pressure on the welfare state (Hatton, 2016; Cavaille, 2017). A number of studies find that when welfare states are generous, immigrants are more likely to be seen as a threat to welfare state generosity, especially when those immigrants are perceived to be culturally or ethnically different from the host

[3]Of course, ethnic cleavages do exist in Latin America. For example, many individuals throughout Spanish-speaking Latin America do not speak Spanish and speak an indigenous language instead. However, it is implied from Rodrik's work that these cleavages are of comparatively smaller scale than those in the US and Europe.

country (Alesina et al., 2001; Alesina and Tabellini, 2004). As such, I expect to find evidence for this in my analysis.

Hypothesis 2: *Globalization shocks in the form of immigration will increase support for right-populist parties.*

Hypothesis 2a: *In cases of immigration shocks, support for right-populist parties is moderated by the size and generosity of the welfare state, with larger welfare states producing more support for right-populists.*

Research suggests trade shocks (i.e., job losses due to import competition) from globalization end up benefiting right-populists. Broz et al. (2021) find that communities who saw the largest economic declines as a result of free trade were more likely to vote for right-populist Donald Trump in 2016. This corroborates much of what has been found in the China Shock literature. For example, Colantone and Stanig (2018) found that concerns about immigration became salient only in regions hit hard by import shocks, further suggesting that trade shocks lend support to right-populists as opposed to left populists.

Hypothesis 3: *Globalization in the form of trade shocks will increase support for right-populists.*

Hypothesis 3a: *This effect will interact with the size of the welfare state & immigration levels in the expected directions (more immigration, more support for right populists, especially with a larger welfare state)*

An important variable to consider in modeling support for populists is whether a country faces a public debt crisis. Hopkin (2020) highlights the need to control for the creditor or debtor status of a country, as he finds evidence that right-wing anti-system parties (an umbrella term under which populist parties fall) in Europe tend to be more successful in creditor countries while debtor countries tend to support anti-system left-wing parties because of their support for anti-austerity measures that negatively impact the breadth the welfare state. Greece's Syriza government is the paradigmatic example of this kind of politics in the 21st century. Like Hopkin, I argue that debtor status is more amenable to left-populist

success than that of right-populists.

Hypothesis 4: Support for left and right populists will be moderated by the Creditor or Debtor status of a country, with Creditor countries increasing support for right-populists and debtor countries increasing support for left-populists.

Research Design

To test these hypotheses, I conducted an exploratory cross-national OLS regression analysis in which the dependent variable is the vote share of left-populist and right-populist parties. Key independent variables included changes in migrant stock, banking crises, trade shocks from China, welfare state universality, and ethnic fractionalization.

Electoral success was measured using the Global Party Survey's expert-coded classification of parties as left or right populist in their rhetoric and policy, as well as each party's electoral (Norris, 2020). Parties are classified as:

1: Pluralist–Liberal.

2: Pluralist–Conservative.

3: Populist–Liberal.

4: Populist–Conservative.

Pluralist parties are considered by Norris (2020) to represent the opposite of populist politics. Thus, this variable quickly tells us both if a party is 1) populist and 2) on the ideological left or right. For example, UKIP in the UK would be classified as "4" and their vote share would be measured by the following variable.

Independent Variables

Welfare State Generosity and Universality: An often preferred metric for welfare state generosity is the post-tax benefit replacement rate for an individual, also known as the Unemployment Insurance Replacement Rate (UIRR), a continuous, numerical value between

0 and 1 measuring the proportion of a person's income that is replaced by unemployment insurance (Hays et al., 2005; Scruggs et al., 2017). However, data on UIRR is only available across 33 countries, while the Global Party Survey offers data for 163 countries. As such, I opt for the use of welfare state universality as measured by V-Dem, which describes the extent to which welfare benefits are available to all adults regardless of their income (Coppedge et al., 2022). I consider welfare state universality an important dimension through which one can understand the breadth of a welfare state, as it offers a sense of a country's credible commitment to the widespread provision of benefits. In future iterations of this study, a comprehensive global dataset of UIRR should be constructed to couple with V-Dem's universality measure.

Net Change in Migrant Stock: Immigration levels are measured as a percentage of a country that is foreign born using the United Nations' yearly updated migrant stock data. Net Change in Migrant Stock is be calculated as the percent difference in migrant stock since the previous year. In regional analysis scholarship prefers the shift-share model of immigration, which accounts for the fact that immigrants tend to move to places where other immigrants already are, and thus measures the shift in the share of immigrants in a particular community (Alesina and Tabellini, 2004). However, this is a subnational measure of immigration, and I intend to use country levels of immigration to do cross-national analysis.

One of the benefits of a national level migrant stock approach is that it exploits how voters often only view immigration as politically salient when it has become a *nationally* politicized issue via media framing. In this way, communities with very few immigrants can still see big changes in immigration attitudes because of a large influx in a different part of the country (Hopkins, 2010).

Ethnic Fractionalization: I measure ethnic fractionalization using Drazanova (2020)'s Historical Index of Ethnic Fractionalization which describes the probability that two randomly selected individuals in a given country will hail from a different ethnic group. For

example, this index has a value of .22 for the US in 1945 (the first year for which there is data), and a value of .53 for the US in 2013, meaning the chance of two randomly selected individuals being from a different ethnic group in the US increased from 22% to 53% from 1945-2013.

Economic Shocks: Globalization shocks in the form of rapid increases of Chinese imports are measured via World Bank data as percent change in Chinese imports as a percentage of GDP. I also include the labor force participation rate as a proxy for economic distress, as measured by the World Bank. To measure financial crises, I use Laeven and Valencia (2020)'s IMF Banking Crises Database II, which provides data on systemic banking crises along with other useful variables like the fiscal cost of a crisis as a percentage of GDP or of public debt.

Results

I show the results of four model specifications for both left and right populist electoral success, but first note two concerns with the data. First, I do not have consistent data on Chinese import shocks for most countries in the dataset. As such, I run one model that includes the Chinese imports as a percentage of GDP and one that does not. Secondly, the mechanisms by which I expect economic shocks to produce support for populism are democratic in nature, but not all countries in the Global Party Survey are consolidated democracies. As such, I show model specifications where only those countries that achieve the status of electoral democracy according to V-Dem's Regimes of The World classification are included. Competitive authoritarian states, for example, are not included in these analyses. Given the exploratory nature of these analyses and the number of methodological concerns, these results should be considered preliminary and used as the basis for outlining future research.

I begin with presenting the correlates of right-populist electoral success in Table 1. When

all elections in the Global Party Survey are included in the analysis (N = 132), the only variable associated with increased support for right-populists is an increase in migrant stock. This finding does not hold when the model is reduced to the 35 elections for which data on Chinese import changes are available. In the model which only includes elections in unambiguous electoral democracies, the only statistically significant result is that welfare state universality is negatively correlated with support for right-populists. In no model specifications is the interaction between welfare state universality and migrant stock change significant, as I expected it to be. Increased welfare state universality's association with less support for right populists is also contrary to expectations, though this result only appears in the model specification of elections in unambiguous electoral democracies (N = 55).

Next, I detail the determinants of support for left populists as seen in Table 2. Only one statistically significant predictor of support for left populists emerges, and that is that decreased ethnocultural fractionalization is associated with support for left populists, but only in the model specification where all elections and Chinese import shocks are included in the model (n =35). This finding is in keeping with expectations, as the lack of ethnocultural fractionalization allows left populists to direct electoral anger toward financial elites as opposed to ethnic minorities and immigrants. However, this finding lacks robustness and should not be seen as conclusive. Theoretically, we should also expect decreases in migrant stock, banking crises, and debtor status to be associated with left populists, but this is not observed in any model specifications.

Table 1: Correlates of Right Populist Vote Share

	All obs	All + China Imports	Democracies	Dem. + China Imports
Constant	−1.730	41.936	51.159	42.648
	(11.794)	(58.861)	(37.297)	(66.347)
Welfare universality	−0.177	−7.804	−18.978**	−20.116
	(1.902)	(12.223)	(5.557)	(10.673)
Change in migrant stock	2.931***	5.517	−1.938	−1.406
	(0.600)	(7.038)	(4.329)	(5.719)
Ethnocultural fractionalization	−4.851	21.530	6.436	19.389
	(8.820)	(18.127)	(18.327)	(18.834)
Banking crisis	8.979	−2.357	−0.706	2.350
	(5.431)	(9.071)	(8.920)	(10.537)
GDP per capita (USD $10,000s)	1.092	1.206	1.529	−0.310
	(1.455)	(3.510)	(2.284)	(4.162)
Labor force participation	0.266	−0.428	−0.292	−0.142
	(0.173)	(1.007)	(0.654)	(1.154)
Creditor	2.271	3.127	7.211	6.315
	(5.551)	(7.844)	(6.779)	(8.065)
Welfare univ. x Migrant change	−2.304	−1.127	1.947	3.841
	(1.168)	(4.349)	(2.269)	(3.662)
China import shock (% GDP)		−0.009		0.051
		(0.293)		(0.311)
N	132	35	55	32
R^2	0.098	0.196	0.177	0.208
Adj. R^2	0.039	−0.094	0.034	−0.115

* p <0.05, ** p <0.01, *** p <0.001
DV: right populist vote share (%). Pooled OLS; SEs clustered by country.
China-shock columns run on reduced sample when observed.

Table 2: Correlates of Left Populist Vote Share

	All obs	All + China Imports	Democracies	Dem. + China Imports
Constant	17.020	33.676	89.716*	47.396
	(14.254)	(44.508)	(43.870)	(53.569)
Welfare universality	0.710	−4.524	−6.851	0.519
	(2.517)	(10.514)	(7.233)	(13.452)
Change in migrant stock	0.932	−1.290	−0.872	0.863
	(0.638)	(3.981)	(4.054)	(4.820)
Ethnocultural fractionalization	−7.525	−52.530*	9.032	−43.633
	(9.914)	(24.089)	(21.539)	(23.624)
Banking crisis	11.711	11.084	−0.388	5.158
	(6.765)	(10.425)	(8.855)	(10.008)
GDP per capita (\$10{,}000s)	−2.492	−2.884	−2.671	−1.291
	(1.818)	(3.729)	(2.382)	(3.769)
Labor force participation	0.048	0.157	−1.029	−0.245
	(0.246)	(0.866)	(0.758)	(0.906)
Creditor	7.426	3.640	8.603	0.547
	(5.838)	(10.134)	(8.033)	(10.647)
Welfare univ. x Migrant change	0.202	1.073	1.845	−0.782
	(1.356)	(2.968)	(2.490)	(3.822)
China import shock (% GDP)		0.207		0.127
		(0.339)		(0.364)
N	132	35	55	32
R^2	0.048	0.259	0.112	0.178
Adj. R^2	−0.014	−0.008	−0.042	−0.158

* p <0.05, ** p <0.01, *** p <0.001
DV: left populist vote share (\%). Pooled OLS; SEs clustered by country.
China-shock columns run on the reduced sample when observed.

Discussion

The right-populist wave of the 2010's and 2020's has brought a great deal of scholarly interest to the determinants of support for candidates and parties whose anti-elite messaging is often coupled with anti-immigrant and hyper-nationalist politics. Yet, left-populists are a formidable political force that can also destabilize democracies. In fact, one of the earliest recorded descriptions of a democracy descending into authoritarianism is that of a left-populist demagogue in Plato's *Republic*, while more recent examples include the authoritarian regime of Hugo Chavez of Venezuela and the democratic Syriza government in Greece. Of late, little research has focused on why electorates opt for right-populists as opposed to left-populists when they are in the political mood to reject the mainstream. This is a particularly important question because it helps explain why an electorate might opt for a President Donald Trump as opposed to a President Bernie Sanders.

Rodrik (2021), I believe, rightly argues that welfare state contexts are key to explaining why voters sometimes opt for left populists over right populists. In particular, strong welfare states coupled with strong ethnocultural cleavages and large levels of immigration create conditions that ethnonationalist right-populists readily exploit. In particular, the phenomenon of welfare chauvinism, where right populists argue that immigrants or minorities are less deserving of welfare provisions can be a powerful political message. Left-populists, on the other hand, should be more successful in more homogeneous societies facing financial crises, where the ire of the public can be more easily directed toward a global financial elite as candidates promise to reject austerity measures. Yet, little empirical work has been done to differentially test the determinants of support for left-populism as opposed to right-populism. As such, I construct and test a modified version of Rodrik (2021) s theoretical framework.

The findings I present are exploratory, and it is for this reason that I highlight ways to strengthen causal inference for future iterations of this study. My main finding is that, in some model specifications, more ethnically and culturally homogeneous societies show higher support for left-populists, while increases in migrant stock heighten support for right-

populists. These findings are in keeping with expectations, however, no significant interactions with the universality of the welfare state are found. Moreover, I was unable to test whether UIRR (a key metric of welfare state generosity) could help explain why more homogeneous societies saw more support for left populist parties in the elections analyzed.

I caution strongly against drawing conclusions from these results and highlight a number of ways to improve this study's methodology. First, globalization related shocks need to be better modeled. This can be done by comprehensively modeling large changes in *all* imports and exports, not just those from China, to better capture the political economy of trade. Other globalization related financial shocks can be measured as well, as Broz et al. (2021) effectively does. These include manufacturing employment share – of particular relevance to de-industrialized Western democracies. Additionally, welfare state generosity as measured by UIRR is not included in this study, and would offer another important dimension of analysis to a story where welfare state contexts are key. Additionally, this study measures banking crises as a binary variable, but a more granular measure of the severity of said crises would allow for more precise modeling. Rodrik (2021) also highlights the importance of *globalization* related financial shocks in explaining support for populists, but the Laeven and Valencia (2020) Systemic Banking Crises Database II does not categorize shocks as being globalization related. Future research should manually code shocks by type for more precise causal inference.

The welfare state plays a unique role in populist politics. Left-populism and right-populism share similar roots: both argue that a virtuous people is exploited by a corrupt elite (Norris and Inglehart, 2019). Unlike traditional left-right party divides, both left and right populist parties of both varieties are typically protective of the welfare state (Hopkin, 2020). Theoretical expectations indicate that welfare state contexts should shape whether voters support left-populists who argue for welfare state expansion to the chagrin of financial elites, or right-populists who view immigrants and minorities as posing an undue burden on the welfare state. This study offers extremely preliminary empirical insight into the relationship

between the welfare state, immigration, and differential support for left and right populists. Future research that builds upon this will offer policymakers and analysts a clearer picture of why populist political backlashes take the shape that they do.

References

Daron Acemoglu and Pascual Restrepo. Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244, June 2020. ISSN 0022-3808. doi: 10.1086/705716. URL https://www.journals.uchicago.edu/doi/10.1086/705716.

Alberto Alesina, Edward Glaeser, and Bruce Sacerdote. Why Doesn't the US Have a European-Style Welfare System?, October 2001. URL https://www.nber.org/papers/w8524.

Alberto F. Alesina and Guido Tabellini. Bureaucrats or Politicians?, February 2004. URL https://papers.ssrn.com/abstract=415046.

Paris Aslanidis and Cristóbal Rovira Kaltwasser. Dealing with populists in government: the SYRIZA-ANEL coalition in Greece. *Democratization*, 23(6):1077–1091, September 2016. ISSN 1351-0347. doi: 10.1080/13510347.2016.1154842. URL https://doi.org/10.1080/13510347.2016.1154842. _eprint: https://doi.org/10.1080/13510347.2016.1154842.

David H. Autor, David Dorn, and Gordon H. Hanson. The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade. *Annual Review of Economics*, 8 (Volume 8, 2016):205–240, October 2016. ISSN 1941-1383, 1941-1391. doi: 10.1146/annurev-economics-080315-015041. URL https://www.annualreviews.org/content/journals/10.1146/annurev-economics-080315-015041.

J. Lawrence Broz, Jeffry Frieden, and Stephen Weymouth. Populism in Place: The Economic Geography of the Globalization Backlash. *International Organization*, 75(2):464–494, February 2021. ISSN 0020-

8183, 1531-5088. doi: 10.1017/S0020818320000314. URL https://www.cambridge.org/core/journals/international-organization/article/populism-in-place-the-economic-geography-of-the-globalization-backlash/98ED873D925E0590CB9A78AEC68BB439.

Charlotte Cavaille. How Distributional Conflict over In-Kind Benefits Generates Support for Far-Right Parties | The Journal of Politics: Vol 85, No 1, 2017. URL https://www.journals.uchicago.edu/doi/abs/10.1086/720643.

Italo Colantone and Piero Stanig. Global Competition and Brexit. *American Political Science Review*, 112(2):201–218, May 2018. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055417000685. URL https://www.cambridge.org/core/journals/american-political-science-review/article/global-competition-and-brexit/C843990101DB9232B654E77130F88398.

Michael Coppedge, John Gerring, and Carl Henrik. Dataset Archive – V-Dem, 2022. URL https://www.v-dem.net/data/dataset-archive/.

Rafael Di Tella and Dani Rodrik. Labour Market Shocks and the Demand for Trade Protection: Evidence from Online Surveys. *The Economic Journal*, 130(628):1008–1030, May 2020. ISSN 0013-0133. doi: 10.1093/ej/ueaa006. URL https://doi.org/10.1093/ej/ueaa006.

Lenka Drazanova. Historical Index of Ethnic Fractionalization Dataset (HIEF), September 2020. URL https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/4JQRCL.

Menno Fenger. The social policy agendas of populist radical right parties in comparative perspective. *Journal of International and Comparative Social Policy*, 34(3):188–209, October 2018. ISSN 2169-9763, 2169-978X. doi: 10.1080/21699763.2018.1483255.

Manuel Funke, Moritz Schularick, and Christoph Trebesch. Going to extremes: Politics after financial crises, 1870–2014. *European Economic Review*, 88:227–260, September 2016. ISSN 0014-2921. doi: 10.1016/j.euroecorev.2016.03.006. URL https://www.sciencedirect.com/science/article/pii/S0014292116300587.

Wiliam Galston. The Populist Challenge to Liberal Democracy, 2018. URL https://www.journalofdemocracy.org/articles/the-populist-challenge-to-liberal-democracy/.

Gordon Hanson, Chen Liu, and Craig McIntosh. The Rise and Fall of U.S. Low-Skilled Immigration. *Brookings Papers on Economic Activity*, 2017(1):83–168, 2017. ISSN 1533-4465. URL https://muse.jhu.edu/pub/11/article/671748.

Timothy J. Hatton. Immigration, public opinion and the recession in Europe. *Economic Policy*, 31(86):205–246, April 2016. ISSN 0266-4658. doi: 10.1093/epolic/eiw004. URL https://doi.org/10.1093/epolic/eiw004.

Jude C. Hays, Sean D. Ehrlich, and Clint Peinhardt. Government Spending and Public Support for Trade in the OECD: An Empirical Test of the Embedded Liberalism Thesis. *International Organization*, 59(2):473–494, April 2005. ISSN 1531-5088, 0020-8183. doi: 10.1017/S0020818305050150.

Jonathan Hopkin. Introduction: A Quick History of the Present. In Jonathan Hopkin, editor, *Anti-System Politics: The Crisis of Market Liberalism in Rich Democracies*, page 0. Oxford University Press, June 2020. ISBN 978-0-19-069976-5. doi: 10.1093/oso/9780190699765.003.0001. URL https://doi.org/10.1093/oso/9780190699765.003.0001.

Daniel J. Hopkins. Politicized Places: Explaining Where and When Immigrants Provoke Local Opposition. *American Political Science Review*, 104(1):40–60, February 2010. ISSN 1537-5943, 0003-0554. doi: 10.1017/S0003055409990360.

Rhoda E. Howard-Hassmann. The Right to Food Under Hugo Chávez. *Human Rights Quarterly*, 37(4):1024–1045, 2015. ISSN 1085-794X. URL https://muse.jhu.edu/pub/1/article/603225.

Miriam Kornblith. Latin America's Authoritarian Drift: Chavismo After Chávez? *Journal of Democracy*, 24(3):47–61, 2013. ISSN 1086-3214. URL https://muse.jhu.edu/pub/1/article/512752.

Luc Laeven and Fabian Valencia. Systemic Banking Crises Database II. *IMF Economic Review*, 68(2):307–361, June 2020. ISSN 2041-417X. doi: 10.1057/s41308-020-00107-3. URL https://doi.org/10.1057/s41308-020-00107-3.

Cas Mudde and Cristóbal Rovira Kaltwasser. Populism. In Michael Freeden and Marc Stears, editors, *The Oxford Handbook of Political Ideologies*, page 0. Oxford University Press, August 2013. ISBN 978-0-19-958597-7. doi: 10.1093/oxfordhb/9780199585977.013.0026. URL https://doi.org/10.1093/oxfordhb/9780199585977.013.0026.

Pippa Norris. Global Party Survey, 2019, April 2020. URL https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/WMGTNS.

Pippa Norris and Ronald Inglehart. *Cultural Backlash: Trump, Brexit, and Authoritarian Populism*. Cambridge University Press, Cambridge, 2019. ISBN 978-1-108-42607-7. doi: 10.1017/9781108595841. URL https://www.cambridge.org/core/books/cultural-backlash/3C7CB32722C7BB8B19A0FC005CAFD02B.

Reuters. Factbox: Venezuela's nationalizations under Chavez. *Reuters*, October 2012. URL https://www.reuters.com/article/world/factbox-venezuelas-nationalizations-under-chavez-idUSBRE89701X/.

Kenneth M. Roberts. Latin America's Populist Revival. *SAIS Review of International Affairs*, 27(1):3–15, 2007. ISSN 1945-4724. URL https://muse.jhu.edu/pub/1/article/212482.

Dani Rodrik. Why Does Globalization Fuel Populism? Economics, Culture, and the Rise of Right-Wing Populism. *Annual Review of Economics*, 13(Volume 13, 2021):133–170, August 2021. ISSN 1941-1383, 1941-1391. doi: 10.1146/annurev-economics-070220-032416. URL https://www.annualreviews.org/content/journals/10.1146/annurev-economics-070220-032416.

John Gerard Ruggie. International Regimes, Transactions, and Change: Embedded Liberalism in the Postwar Economic Order. *International Organization*, 36(2):379–415, 1982. ISSN 0020-8183. URL https://www.jstor.org/stable/2706527.

Lyle Scruggs, Jahn Detlef, and Kati Kuitto. Comparative Welfare Entitlements Dataset (CWED2), 2017. URL http://cwed2.org/.

"Compensate the AI Losers?" A Video Experiment on AI-Framing and Policy Preferences

Abstract

What impact does information about AI have on policy preferences? Using a pre-registered video survey experiment (N=1,009) of U.S. white-collar workers, I examine how information about AI as either a job threat or productivity tool shapes policy preferences for redistribution (taxes and transfers) and "predistribution" (regulatory policies aimed at reducing inequality before taxes and transfers). In the absence of AI-information, I find voters are pessimistic about AI's wider economic impact, but are cautiously optimistic about their own prospects. However, once exposed to videos about AI, voters are far more responsive to AI-threat frames than optimistic ones. Partisanship and education emerge as key lenses through which voters interpret the AI threat. Republicans respond to AI threat videos by rejecting social safety nets, but no significant results emerge for Democrats. Those with less education prefer predistribution in the form of AI regulation over wealth redistribution. My findings suggest a potential opening for Democrats seeking to regain their dwindling share of the electorate that does not hold a college degree by offering voters AI regulation.

Keywords: artificial intelligence, public policy, survey experiment, public opinion, political economy

Introduction

What does the information voters receive about AI have on their policy preferences? In a 2023 CNBC interview, IBM CEO Arvind Krishna made a statement that alarmed many workers:

> "I do believe that AI is going to replace a lot of what I'm calling 'white-collar clerical jobs' ... I think a good 30 percent of those roles could go away over five years."

However, he quickly added:

> "At the same time, it will create lots of other roles ... I believe our total employment will increase" (CNBC, 2023).

The extent to which AI will ultimately displace workers is currently up for debate, and there is no shortage of conflicting predictions regarding AI's labor market impacts in the media. Popular messages include: AI will cause massive job losses (Sigalos, 2025; Khan, 2025), AI will have little to no impact on the labor market (Milmo, 2025), or that AI will *complement* human workers and help them earn more money (Zahn, 2025). Voters' about AI's economic impact should theoretically shape their policy preferences, as research from the "pre-AI era" suggests that voters should support increased redistribution (Busemeyer and Sahm, 2022; Cusack et al., 2008; Ravallion and Lokshin, 2000) and regulation (Dancygier and Donnelly, 2013; Owen and Johnston, 2017) as insurance against threats to their financial well-being. Why then, does recent research find that information about AI's economic impacts almost always fails to produce changes in preferences for redistribution and regulation (Heinrich and Witko, 2025; Margalit and Raviv, 2023; Zhang, 2022; Menon and Zhang, 2025)?

I suggest three explanations for the literature's puzzling findings. First, prior studies largely rely on text-based vignettes to expose respondents to AI information, which may not properly recreate voters' increasingly common experience of obtaining economic and

political information in online videos (Newman et al., 2025). Second, I posit that individuals' perceptions of AI's broader impact on the economy are pessimistic and sticky, making their policy preferences generally less responsive to information about AI. Finally, partisanship and education level play key, largely under-analyzed roles in shaping voters' preferred AI-related policy preferences. Both partisan identity and degree attainment are strong predictors of economic policy preferences, and should naturally inform how voters see the government's role in addressing challenges from AI, with Democrats being generally more amenable to state intervention to advance economic policy goals than Republicans (Gelman, 2007).[1]

To address economic challenges, research from the pre-AI era shows degree holders tend to favor redistribution (e.g., taxes and transfers to fund unemployment insurance, low-cost healthcare, food stamps) while those without degrees favor what scholars increasingly refer to as predistribution in the form of pre-tax interventions like labor market regulations, minimum wage laws, and job retraining programs (Hacker, 2011; Kuziemko et al., 2023). For example a voter without a degree who is concerned about AI may prefer the government raise the minimum wage in response or simply regulate the technology, but would not favor redistribution in the form of raising taxes on the wealthy to fund unemployment insurance. Measuring policy preferences in response to AI must therefore differentiate between predistribution and redistribution.

To test these claims, I fielded a preregistered video-based survey experiment (N=1009) among white-collar workers in the U.S., randomly assigning respondents to view videos framing AI as either substituting white-collar labor or complementing it. I then measured expectations about AI's impact on the economy and policy preferences for redistribution and predistribution. Respondents are largely unconvinced by optimistic information about AI, while AI-threat information darkens their economic forecasts. Importantly, the videos produce changes in individuals' policy preferences that vary by education level and party. In

[1]It is worth noting that the Republican Party under the leadership of Donald Trump undertakes state interventions in the form of protective tariffs, bailouts for tariff-effected farmers, and even taking strategic stakes in large companies (Morescalchi, 2025)

response to both optimistic and pessimistic information about AI's impact on the white-collar workforce, Republicans prefer fewer labor protectionist policies in the form of job retraining for those who lose their jobs to AI, while no significant effects are found for Democrats. Those without a college degree prefer AI regulation (a form of predistribution) as a policy response to the AI threat, not redistribution.

AI adoption has nearly doubled since 2023, when most experiment-based studies of AI-related policy preferences were conducted (Pendell, 2025). In this context, this study offers scholars and policymakers an updated picture voters' policy preferences using an experimental treatment method that better replicates voters' real life experiences with AI-information: videos (Mutz, 2021; Wittenberg et al., 2021). Voters' beliefs about the labor market impacts of AI may precipitate a meaningful political backlash, and the added external validity of video treatments helps clarify how information about AI shapes policy preferences.

First, this study expands the literature on predistribution into the AI-era. While previous experiment-based studies make important contributions to the political economy of AI, this study is the first to show that the less educated (and often most vulnerable to AI) prefer predistribution in the form of AI regulation over redistribution as a response to negative AI frames (Heinrich and Witko, 2025; Golin and Rauh, 2022; Magistro et al., 2025). These findings underline the continued relevance of the predistribution vs. redistribution model and carry political implications with them (Hacker, 2011; Kuziemko et al., 2023; Bozio et al., 2024). As policymakers debate AI regulation, my findings suggest that Democrats may regain their dwindling share of voters without a degree with calls to regulate the technology.

Moreover, I contribute to this literature by shedding light on the emerging partisan dynamics of AI-related policy preferences. In early 2025, AI began to take on limited partisan valence as leaders of AI firms began uneasy alliances with the Trump Administration. However, the vast majority of studies seeking to understand AI-related policy preferences in the US were conducted before 2025 and are unable to speak to this dynamic (Golin and Rauh, 2022; Green et al., 2025; Vesely and Amaris, 2025). Expanding on this research, I show that

Republicans respond to the labor market disruptions of AI differently than Democrats, preferring to reduce social safety nets (in the form of job retraining and UBI) when faced with both AI-threat and AI-complementarity frames. These results build upon those of Magistro et al. (2025), whose 2023 survey experiment showed an alternative partisan dynamic, where Republicans responded to information about AI driving up inflation with increased support for job retraining. The present study's more recent results suggest the beginnings of partisan shift in the perception of AI policy in the US at the beginning of Trump's second term.[2]

Addressing AI's labor market impacts is a policy platform yet to be fully captured by either major U.S. party (Gallego and Kurer, 2022). Such a lack of political ownership represents a potential electoral boon to the party that can harness it. An instructive recent example of this kind of political maneuvering can be found in the Trump campaign's strategy of appealing to workers displaced by globalized trade. In many ways, the political economy of AI resembles that of globalization just a few decades ago, where white-collar workers today are analogous to manufacturing workers [3] (Magistro et al., 2025). By treating voters with videos highlighting AI's capabilities to transform the white collar labor market, the present study indicates that AI may represent a similarly untapped political opportunity. Voters who do not hold a degree have proven instrumental in deciding recent elections in the US, and this study shows that offering predistribution in the form of AI regulation may be a successful strategy in courting this influential bloc (Goidel et al., 2024).

Technology and Policy Preferences

Information about automation should not have identical impacts on voters because automation can threaten some workers by cheaply replacing the tasks they complete, while offering

[2] I state this cautiously as their sample was not limited to white-collar workers (as was the present study) and highlighted AI-induced price shocks, not job threats.

[3] Like support for AI investment, support for globalization in the form of free trade agreements like NAFTA was largely bipartisan, and an influential subset of the electorate in key swing states was disproportionately negatively impacted by the ensuing labor market changes. This is certainly not the full story of the 2016 election, but Donald Trump brought these anti-globalization voters into his coalition and catalyzed tremendous changes in both US and global politics (Margalit, 2019).

others potential wage and productivity gains (Acemoglu and Johnson, 2023; Loaiza and Rigobon, 2024). AI[4] is no different than previous waves of automation in this regard, however, it stands to impact white-collar workers, a subset of the labor market that has until recently stood to benefit tremendously from advances in information technology (Autor et al., 2003; Felten et al., 2021). These workers' expectations of AI's economic impact are extremely politically relevant because policy attitudes are often driven by one's expected future income (EFI) (Checchi and Filippin, 2003; Ravallion and Lokshin, 2000) and sociotropic concerns about the wider economy (Fong, 2001). Moreover, workers' economic forecasts are often influenced by their information environments (Binder et al., 2025; Druckman et al., 2012; Garz, 2024). Thus, information about AI's labor market impacts should impact economic policy preferences. Given that white-collar workers make up roughly 60% of the workforce in the US, this group's expectations of AI's impact on the economy and preferred policy responses can have significant electoral impacts (Staiti, 2024).

When workers learn that a new technology will reduce their income, increase their likelihood of unemployment, or negatively affect the economy as a whole, they often demand economic policy responses. Broadly speaking, there are two ways policymakers can respond to AI-related job concerns. First, they can invest more in redistribution, typically conceived of as taxes and transfers to fund programs like unemployment insurance, low-cost healthcare, and food stamps. Alternatively, they can employ predistribution, which refers to policies that aim to alleviate inequality *before* taxes and transfers. In his influential work, Hacker

[4]It is debatable whether Artificial Intelligence is the appropriate term to use to describe this technology. Services such as ChatGPT, Claude, and Gemini are primarily Large Language Models (LLMs). LLMs are extremely sophisticated machine learning (ML) based prediction programs trained on vast amounts of data. By identifying patterns in human-written text, for example, and being trained through processes like Reinforcement Learning from Human Feedback (RLHF) where a human ranks and evaluates LLM responses to queries, these models' text prediction has become so precise that they typically offer competent responses to complex questions.

Whether or not LLMs should be referred to as "intelligent" is contentious considering that these programs do not reason in the way that humans do. An obvious example of their somewhat alien process of producing text is in the problem of "hallucination", where an LLM will create false information that *sounds* like it could be correct based on its text-prediction model and report it to the user in a confident tone. For the purposes of this paper, I refer to LLMs like ChatGPT, Claude, and Gemini as AI because that is how they are discussed in common parlance.

(2011) describes predistribution as a "focus on market reforms that encourage a more equal distribution of economic power and rewards even before government collects taxes or pays out benefits." Examples of predistribution include but are not limited to labor market regulations, jobs guarantees or job training programs, and unionization protections. A growing literature highlights the distinction between redistribution and predistribution because different groups of voters often favor one over the other in response to the same economic challenges (Bozio et al., 2024; Hacker, 2011; Kuziemko et al., 2023). A logical implication of these findings is that conceiving of economic policy preferences as falling on a spectrum of pro or anti redistribution lacks necessary nuance. The present study adopts the predistribution vs. redistribution framework because of its demonstrated explanatory power.

Research from the pre-AI[5] era indicates workers often support regulation—a key form of predistribution—in response to generic labor market threats (Dancygier and Donnelly, 2013; Owen and Johnston, 2017). More recent work suggests these preferences are multidimensional and influenced by partisan identity and education levels, as reflected in the growing proportion of high-income, educated Democrats who support redistribution seemingly against their own narrow economic interests (Kuziemko et al., 2023; Magistro et al., 2025; Vilbig, 2026). Furthermore, workers support increased redistribution as insurance against threats to their expected future incomes and employment, such as those represented by AI (Busemeyer and Sahm, 2022; Ravallion and Lokshin, 2000; Cusack et al., 2008)). Alternatively, those with *upward* income trajectories often want less redistribution in order to protect future earnings (Benabou and Ok, 2001). All in all, the policy literature from the pre-AI era indicates that voters largely respond to economic threats like recessions or trade shocks with support for both increased predistribution and redistribution, while education is a strong predictor of which type they favor.

Thus, one would expect information about the AI job threat to move voters' policy

[5]Debates exist over what exact technologies can be considered AI and thus when the AI era actually began. For the purposes of this paper I consider the pre-AI era to be prior to the release of ChatGPT in 2022, when a large portion of the public gained access to Large Language Models.

preferences. However, prior research on AI-information and policy preferences offers mixed findings on this front. Some studies find negative AI frames increase support for redistribution (Golin and Rauh, 2022), while more recent research finds effects only among specific subgroups like those with high AI knowledge (Heinrich and Witko, 2025), and many find null effects (Zhang, 2022; Margalit and Raviv, 2023). There is a growing consensus that AI-related policy attitudes are motivated by sociotropic concerns (fear about general job loss in your community or country, as Haslberger et al. (2025) and Zhang (2022); though some evidence points to the importance of egoistic concerns (personal financial worries, as Golin and Rauh (2022) suggest), or both sociotropic and egoistic concerns, as argued by Heinrich and Witko (2025).

Nearly all work on this subject uses text-based AI frames, which respondents may not considered as believable as videos (Wittenberg et al., 2021). Additionally, most AI-specific survey experiments were fielded in 2023 or earlier, before widespread AI adoption. Many were also fielded outside of the US. Haslberger et al. (2025), for example, conducted a study in the United Kingdom that did not rely on text-based vignettes and instead asked respondents to use ChatGPT. The authors found that after having respondents treated respondents support increased predistribution in the form of jobs guarantees, but not redistribution, indicating that voters do indeed distinguish between these policy types. In the US, a comparatively weaker welfare state may make respondents more sensitive to the AI threat and thus more willing to shift policy preferences in the face of information (Busemeyer and Sahm, 2022; Haslberger et al., 2025; Green, 2023; Vesely and Amaris, 2025).

If one extends the policy logic of the pre-AI era to today, white-collar workers should support both increased redistribution and predistribution in response to information highlighting the AI threat. An example of this kind of framing is that of Meta CEO, Mark Zuckerberg, who in stated AI will replace mid-level software engineers on the Joe Rogan Experience podcast (read, 2025). Those shown optimistic information about AI's potential to complement their labor, though, should want less redistribution and predistribution as

a way to protect future economic gains. However, given the amount of research indicating either small or null effects from AI information, I do not expect these policy preference shifts to be large (Haslberger et al., 2025).

H1: Exposure to the Substitution (Complementarity) treatment increases (decreases) support for redistributive and predistributive policies relative to Control.[6]

If voters' policy preferences should, in theory, shift in response to information about economic threats, why does information about AI's disruptive potential so often fail to move them? One reason may be that aggregate analyses often obscure meaningful differences between subsets of the workforce. Consider the case of measuring preference shifts in response to AI-threat information wherein all respondents believe AI will negatively impact the economy. Furthermore, half of respondents think increasing taxes on corporations will help ameliorate AI-driven inequality while the other half thinks *decreasing* taxes on corporations will stimulate the economy and better respond to inequality. While this is an extremely simple case, it shows the potential issues aggregate-level analyses. I argue that partisanship and education level reflect the two most relevant divides in explaining how different individuals respond to information about AI.

Partisanship is the natural starting point for subgroup analyses given the extent to which it predicts policy support. Partisanship informs both the processing of political information and voters' philosophies about the role the state should play in the economy and how it should respond to inequality (Iyengar and Krupenkin, 2018; Bartels, 2006; Taber and Lodge, 2006; Zaller, 1992). Republicans and Democrats may also differ in their personal beliefs regarding the deserving-ness of those who lose jobs due to AI and what the role of government is in ameliorating said job loss. Indeed, previous work has indeed shown that personal beliefs like these impact preferences for redistribution (Bénabou and Tirole, 2006, 2016; Aarøe and

[6]This hypothesis was originally constructed slightly differently, as "Exposure to the Substitution treatment increases support for redistributive and protective policies relative to Control." Furthermore, partisanship was originally considered an exploratory moderator and education was not originally included as a main subset of interest. As such, H2 builds on an exploratory hypothesis and H3 is a post-preregistration hypothesis, but both are strongly theoretically motivated.

Petersen, 2014). I expect that partisans will differ on how to address AI-induced economic impacts.

One existing study gives us some insight into the way partisanship interacts with AI-related policy preferences. A 2023 conjoint experiment showed that partisans in both Canada and the US react differently to messages about AI-related price increases (Magistro et al., 2025). The authors find Republicans' AI-policy attitudes are more malleable in response to information treatments than that of Democrats. In terms of baseline expectations about AI, they find Democrats to have more AI optimism. However, there is reason to believe these dynamics have shifted in light of recent political developments. Since early 2025, a number of top AI firms have entered strategic partnerships with the US government under the leadership of the Trump Administration. The most visible of these is the uneasy alliance between Donald Trump and Elon Musk, who leads several prominent AI and robotics companies.

Other ongoing political realignments further highlight the need for new research. It has historically been the case that Republicans are on more skeptical of state economic intervention in both the form of predistribution or redistribution than Democrats, often because of the concern that its distortionary market effects will worsen the economy while Democrats tend to favor a more active role of the state in alleviating economic disparities (Gelman, 2007). Under the leadership of Donald Trump, however, the Republican party of has increasingly favored state intervention in the economy, as exemplified by the US government taking an $8.9 billion stake in Intel (Morescalchi, 2025), predistribution in the form of increasing tariff rates, and some selective examples of redistribution such as farmer bailouts (USDA, 2025). The Democratic party has also realigned, though over a longer time horizon. Since the 1980's, the Democrats have increasingly preferred redistribution over predistribution (Kuziemko et al., 2023) and has increasingly added high-income, educated voters to its coalition (Vilbig, 2026).

H2: Republicans(Democrats) will be less(more) supportive of both predistribution and redistribution as a response to AI-threats.

College degree attainment is a key dividing line that informs policy preferences and is deeply intertwined with the labor market impacts of AI. While many white collar workers with college degrees saw wage gains from automation in the form of computerization in the late 20th century, much information voters receive about AI centers on the ability of it to automate the labor of these very same educated, white-collar workers who earn their living by operating a computer (Autor et al., 2003; Felten et al., 2021; Schendstok and Wertz, 2024). College degree attainment also predicts a number of relevant characteristics related to AI-policy preferences including partisanship (Zingher, 2022), preferences for predistribution or redistribution (Kuziemko et al., 2023), and, to a lesser extent, vulnerability to automation (Golin and Rauh, 2022; Acemoglu and Restrepo, 2020).

In the pre-AI era, metrics of automation vulnerability largely find that those with Bachelor's degrees or higher are both less concerned about automation and less vulnerable to it (Golin and Rauh, 2022; Acemoglu and Restrepo, 2020). With the rapid proliferation of Large Language Models like ChatGPT and Claude, however, a number of scholars have created new occupational metrics of AI exposure and vulnerability. In these metrics, AI vulnerability is typically understood to mean the actual labor market risk of unemployment or wage decreases due to AI's ability to replace that job. AI exposure, on the other hand, is the amount AI is thought to impact the way a job is done. As one can imagine, AI exposure does not necessarily mean a job is at risk. In some cases, the ability of AI to perform many job tasks can make the worker more productive and potentially be able to command higher wages.

An example of a highly AI exposed job that is not AI vulnerable would be that of a lawyer, who could automate much of their research and email-writing process without fear of job loss. On the other hand, an advertising copywriter may hold a college degree, is similarly highly exposed to AI, and is also extremely vulnerable to being fully replaced by AI. These metrics generally find that the jobs held by those with college degrees are more exposed to AI, though they differ and are often agnostic on the extent to which they indicate degree

holders are *vulnerable*. Part of the reason for this unevenness of vulnerability is because degree holders in some fields, like healthcare and laboratory science are much less vulnerable to degree holders whose jobs can primarily be completed on a computer (Schendstok and Wertz, 2024).

Education does, however, reliably predict preferences for predistribution and redistribution. An empirical analysis of 70 years of American voter data conducted by Kuziemko et al. (2023) shows that predistribution tends to be the policy remedy of choice for the less educated. These dynamics have played out in partisan politics through what is often called the "diploma-divide" wherein Bachelor's degree attainment has become one of the best predictors of identification with the Democratic party since 2000 (Zingher, 2022; Green, 2023). As the degree-holding contingent of the Democratic coalition has grown, so has its favor for redistribution over predistribution as a preferred economic policy choice (Kuziemko et al., 2023).

Individuals also attach economic value to their personal beliefs which may help explain their preferred economic policies (Bénabou and Tirole, 2016), such as those about the importance of not being "lazy" or the notion that people generally "get what they deserve and deserve what they get" (Lerner, 1980). These beliefs can help explain why someone may prefer a regulation that would prevent her from being unemployed by AI in the first place: she places a psychological value on employment that is not captured by a model that views individuals solely as post-tax income maximizers. These beliefs may vary by degree-attainment, which could explain why those without degrees prefer predistribution. [7]

[7]Political economists have not fully investigated *why* the less-educated appear to be more moved by these personal beliefs than those with degrees.

Some research suggests that individuals prefer predistribution because it preserves their social status in a way that redistribution does not. The following anecdote—and a large body of research—is illustrative of this point: post-tax-income maximizing economic models suggest that a receptionist whose job was replaced by AI will prefer an unemployment insurance check (redistribution) that replaced his lost income and allowed him not to work over a regulation meant to prevent their company from using AI to replace them. However, a great deal of research suggests that individuals place value on their pre-tax income that is often tied to personal beliefs like that in the social status of being employed.

For the less educated, it is possible that their increased likelihood of actually needing to use welfare benefits means that the status payoff of predistribution instead of redistribution is more tangible for them than it is for degree holders. In other words, degree holders may be happy to support redistribution over

H3: In response to AI threat information, those without a college degree will prefer increased predistribution.[8]

Research Design

Survey Experiment

To understand how voters' preferences shift in response to AI-information, I fielded a pre-registered video survey experiment (N=1009).[9] White collar workers are the prime target of AI related information, and therefore this study, because of their unique vulnerability to the technology (Felten et al., 2021). Survey respondents were recruited between August 25 and September 7, 2025 through Prolific, a widely used web-survey platform. The survey was completely voluntary, respondents gave informed consent, and its methods were evaluated to meet experimental ethics requirements by IRB. Respondents completed a short Qualtrics survey experiment with a median response time of 8 minutes and 31 seconds. Eligibility was restricted to white-collar workers in the U.S. aged 25–54. Screening of workers was conducted through Prolific, and checked for accuracy via self-reports of occupation. 41.5% of respondents identified as Democrats, 18.3% as Independents, 4.2% as Other, and 36% as Republicans. Respondent ages ranged from 25 to 54, with the average age of a respondent being 39.5. The average income of respondents was $4230/month, or approximately $51,000 annually, and 69% of respondents had a Bachelor's Degree or higher. The final sample of respondents closely reflects the demographic composition of the U.S. white-collar labor force.[10]

predistribution because they are much less likely to actually be recipients of redistribution. Future research should investigate the microfoundations of this difference.

[8]The original preregistration included partisanship, but not education level as a subgroup of interest. This is, however, hypothesis motivated by the large literature on education as a predictor of political attitudes.

[9]The full survey text and video transcripts are available in the Online Appendix

[10]To calculate the partisan, gender, age, and demographic breakdown of the U.S. white-collar workforce, I used public data from the American National Election Survey (ANES), General Social Survey (GSS), and

Most survey experiments in political science utilize text vignettes, however, voters increasingly obtain political and economic information in video format, which research indicates is viewed as being more believable (Mutz, 2021; Wittenberg et al., 2021). Recent studies in political science begun using video based treatments instead of text-vignettes for this reason (e.g., Noordzij et al. (2024)). To better replicate a real-world information environment and thereby increase external validity, the survey employed short video treatments. Respondents were randomly assigned to view one of three short videos with AI information: a video about AI substituting jobs, a video about AI helping white collar workers earning more money, or a control video about a non-AI related topic. I then measured preferences on redistribution and predistribution.

All videos were embedded in the Qualtrics survey and could not be skipped. A post-video attention check confirmed comprehension of the video's message, and participants who failed either of two attention checks were excluded. Of the 1009 respondents, 12 failed at least one attention check, and another 35 were excluded from analysis because their occupational selection indicated they weren't in fact a white-collar worker. After informed consent, participants were randomly assigned to one of three conditions:

Control Group: Randomly shown one of four informational videos that are not related to Artificial Intelligence. These videos were a National Geographic video on giraffes, a video about how to juggle, a clip from a documentary about the Giza pyramids, and a documentary clip about the Grand Canyon.

Complementarity (AI-Optimistic Frame): Shown a video highlighting how AI can save them time and make more money. The video was preceded by the following blurb:

"Several studies published in 2023–2025 report that workers who adopt AI tools can complete tasks faster and take on higher-value work. These analyses estimate potential pay increases of 20–30% for white-collar workers, alongside growth in overall employment in AI-exposed occupations.

Bureau of Labor Statistics (BLS). Demographic tables of the US white-collar workforce are available in appendix table A.1.

Please watch this short video showing how Artificial Intelligence (AI) tools like ChatGPT are helping white collar workers earn more money by helping them complete their tasks faster.

Once you finish watching the video, you will be able to go to the next page."

Substitution (AI-Threat Frame): This group watched a video showing how AI was already leading to layoffs and could do a great deal of white collar work. The video and was preceded by this blurb highlighting the AI-white-collar threat:

"Several studies from 2023–2025 warn that companies may use AI to automate a substantial share of routine professional work. These analyses estimate that about 30% of white-collar workers could face wage reductions or job loss if firms replace tasks previously done by people with AI systems.

Please watch this short informational video showing how Artificial Intelligence (AI) platforms like ChatGPT, Claude, and Gemini are already causing layoffs and reducing wages for white collar workers because of AI's ability to do work cheaply and quickly.

Once you finish watching the video, you will be able to go to the next page."

Outcomes

Immediately following the video, respondents completed a battery of outcome measures. The primary dependent variables were preferences for redistribution and predistribution. To measure redistribution, I created an index of support for questions on raising taxes on the wealthy and using those funds to fund unemployment insurance, healthcare, and social services. Predistribution was measured via support for a variety of regulations on the use of Artificial Intelligence. Regulation of automation is a key form of predistribution because it can prevent job losses by slowing or restrict automation's integration into the economy. Additionally, I measured support for funding the job retraining of workers affected by AI-layoffs, which I consider a case of predistribution in the sense that it alleviates inequality by getting individuals back to work as opposed to sending them direct taxes and transfers. However, it is worth noting that this is not the purest form of predistribution because it

occurs *after* the person loses their job as opposed to preventing them from losing it. Rather, it prevents them for losing a job in the future. Still, job retraining is conceptually distinct from purely redistributive policies that rectify inequality through direct transfers of money like unemployment insurance or healthcare subsidies.

1. ***Redistribution*** [11]

    (a) **Redistribution index** comprised of

        i. General redistribution support[12]

        ii. Taxing the rich

        iii. Taxing corporations

2. ***Predistribution:***

    (a) **Regulation Index** comprised of

        i. Banning AI layoffs

        ii. Mandating AI safety measures

        iii. Creating a federal agency to regulate AI

        iv. General federal AI regulations

Finally, post-treatment manipulation checks and mediator variables captured beliefs about personal and sociotropic economic risks from AI in order to capture the mechanisms through which voters' AI-related policy preferences shift.

[11] I controlled for a number of other post-treatment preferences for universal basic income and union interest, and report results with regard to these variables in the Online Appendix.

[12] The redistribution question read: *"Some people feel that the government should increase taxes a lot and spend much more on health and social services such as unemployment benefits and low-cost healthcare programs like Medicaid. Other people feel that the government should cut taxes a lot and spend much less on health and social services. And other people have views somewhere in-between.Please use the slider to indicate where your views fall on this scale." (0-100 scale)*

Manipulation Checks:

1. Expected income change due to AI (egoistic concern)

2. Perceived AI-related unemployment risk in the next five years (egoistic concern)

3. Expected impact of AI on total number of jobs in the economy (sociotropic concern)

To examine heterogeneity in treatment effects, the survey includes several pre-treatment measures in addition to basic demographics for analysis as independent variables.[13] Particular attention is paid to partisanship and education level. In order to operationalize the potential for a respondent's job to be substituted by AI, I link respondents' self-reported occupations (coded using Standard Occupational Classification (SOC) 2018 codes) to occupational AI risk scores from the EPOCH dataset (Loaiza and Rigobon, 2024).

Analysis

Main results are quantified via t-tests. In most cases, outcome scores are normalized in terms of standard deviations from the control group mean for ease of interpreting effect sizes. I seek treatment effect sizes of at least Cohen's $d = .2$, the conventional threshold for a small effect (Ellis, 2010). For reference, an effect size of .35 would put the power of this study at a similar level to that of Heinrich and Witko (2025), who find effect sizes of 0.34 standard deviations on support for regulating AI when those with high AI knowledge were treated with information about AI's potential to negatively transform the nature of work.

[13]Due to a technical error in importing covariate data for respondents' age and gender from Prolific, 2.5% of respondents do not have observations for age and gender, with full details available in the Online Appendix. In addition to typical demographic variables, I controlled for risk aversion (0–10 scale from Dohmen et al. (2011) and social mobility, captured by asking both one's childhood economic class and one's expected future class. If a person's reported childhood class was "Upper Middle: but their expected future class is "Lower Class", they'd be considered highly downwardly mobile.

Results

Presentation of results begins white-collar workers' preferences on AI-related public policy in the absence of AI-related information. Next, I highlight the impact of video treatments on partisans and the differentially educated. I find that workers' baseline pessimism about AI's impact on the economy is persistent, as positive AI information largely fails to produce optimism about both AI's personal and wider impact. The failure of positive information to produce AI-optimism and of *any* treatment to move policy preferences is contrary to H1's expectations. Second, Republicans react to increased AI-risk by rejecting social safety nets, indicating support for *H2*. Lastly, in support of *H3*, voters with less education respond to AI frames with demands for more AI regulation—a form of predistribution—instead of redistribution.

Baseline AI-Related Public Policy Preferences

Figure 1 shows white-collar workers' economic expectations of AI in the absence of AI information by showing the preferences of the control group, who did not see videos with AI information. I find that white-collar workers are concerned about AI's broader impact on the economy, but not about its impact on them personally. They broadly support increasing spending on redistribution (taxes and transfers) and predistribution (job retraining, AI regulations). Table 1 offers more granular detail, indicating white-collar workers expect AI to have negative impacts on the economy at large by reducing the total number of jobs, they are, on average, unconcerned about the impact AI will have on their own personal employment prospects and *mildly* expect it to increase their future incomes (by 4.77%). Consistent with previous research, sociotropic concerns about AI's economic impact loom larger than personal ones (Haslberger et al., 2025; Heinrich and Witko, 2025). These negative sociotropic expectations dovetail with support for increasing AI regulation. Additionally, white-collar workers are generally supportive of increasing redistribution via spending on social safety

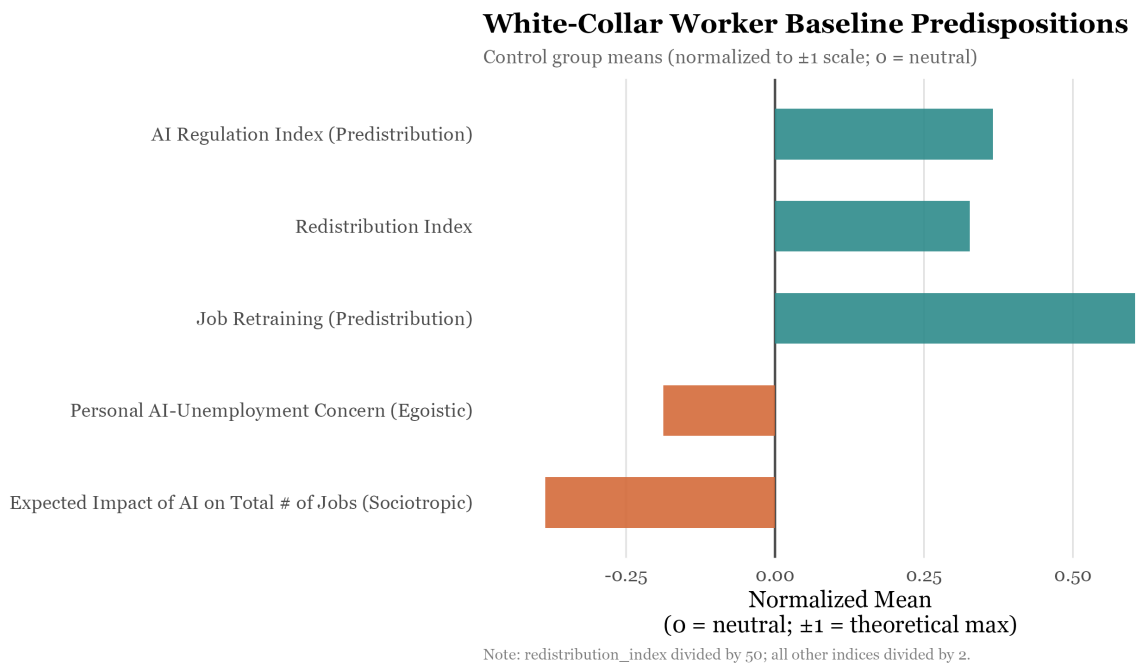


Figure 1: AI-Related Predispositions in Absence of AI information

nets (healthcare spending, unemployment insurance, etc.) funded by increasing taxes on corporations and the rich.

Treatment Effects

I recreated the experience of obtaining AI-optimistic information, labeled as "complementarity", as well as AI-threat information "substitution") in order to determine if AI-information has a meaningful causal impact on voters' policy preferences. As research suggests that automation technologies impact preferences through their impacts on sociotropic and personal/egoistic economic expectations, I measured each treatment's effect on expected future income (egoistic), expectations about the total number of jobs (sociotropic), and perceived personal unemployment risk (egoistic). I estimate the average treatment effect of exposure to AI-threat and AI-positive frames by standardizing the size of the response in standard deviations from the control group mean. For example, the value of .39 on "Regulation Index" for those without a degree exposed to the substitution treatment means they support

AI regulations .39 standard deviations more than their counterparts in the control group. I find that baseline pessimism is sticky as positive frames fail to improve expectations or shift policy preferences. AI threat frames worsen economic forecasts about both themselves and others.

Video treatments framing AI as a labor substitute reduce EFI by 7.4% and makes workers' already pessimistic beliefs about AI's impact on the job market .4 standard deviations worse. Despite these increased fears about AI, I find no treatment effects from either AI information treatment on public policy preferences and fail to reject the null hypothesis with regard to *H1*, though it is worth noting that the substitution treatment is directionally associated with greater support for both predistribution and redistribution. These findings are consistent with Heinrich and Witko (2025) and Zhang (2022), but not Golin and Rauh (2022) and Haslberger et al. (2025).

As noted earlier, there are a number of theoretical reasons to expect that voters' policy responses to AI-threat information would be heterogeneous by nature. Namely, different groups of people prefer different government responses to the same economic threats. For example, Republicans often prefer reducing taxes as a form of stimulating a struggling economy and helping workers, while Democrats often prefer raising taxes to fund redistribution. As shown by Kuziemko et al. (2023), those without a degree tend to prefer predistribution, not redistribution to alleviate inequality. Aggregate level analysis likely obscures how different parts of the electorate react to the AI threat.

Heterogeneous Effects

Given the primacy of partisanship in forming Americans' political attitudes and the increasing politicization of AI via the Trump Administration's partnership with a number of key AI firms, I begin with an analysis of partisans' reaction to information about AI's economic impacts. I then move to results for those with and without college degrees.



Figure 2: Effects of AI-information on Economic Expectations

Partisanship

I find that Republicans respond to both negative and positive information information about AI by rejecting predistribution in the form of job retraining. Republicans also react to AI threat information with less support for Universal Basic Income. No changes in policy attitudes are observed for Democrats. As such, I am partially able to reject the null hypothesis with respect to *H2*. These findings are somewhat consistent with those of Magistro et al. (2025), who report that Republicans primed with AI threats do, like in my sample, support less universal basic income. However, Republicans in the present study react differently in reducing their support for job retraining. One possible explanation for this discrepancy is that individuals respond differently to AI as a consumer price shark than they do to AI as a labor market risk. In the aforementioned authors' study, Republicans faced with an AI shock show increased support for job retraining after taking part in their conjoint experiment which highlights the role of AI in inflating prices (See Magistro et al. (2025) appendix figure G1).

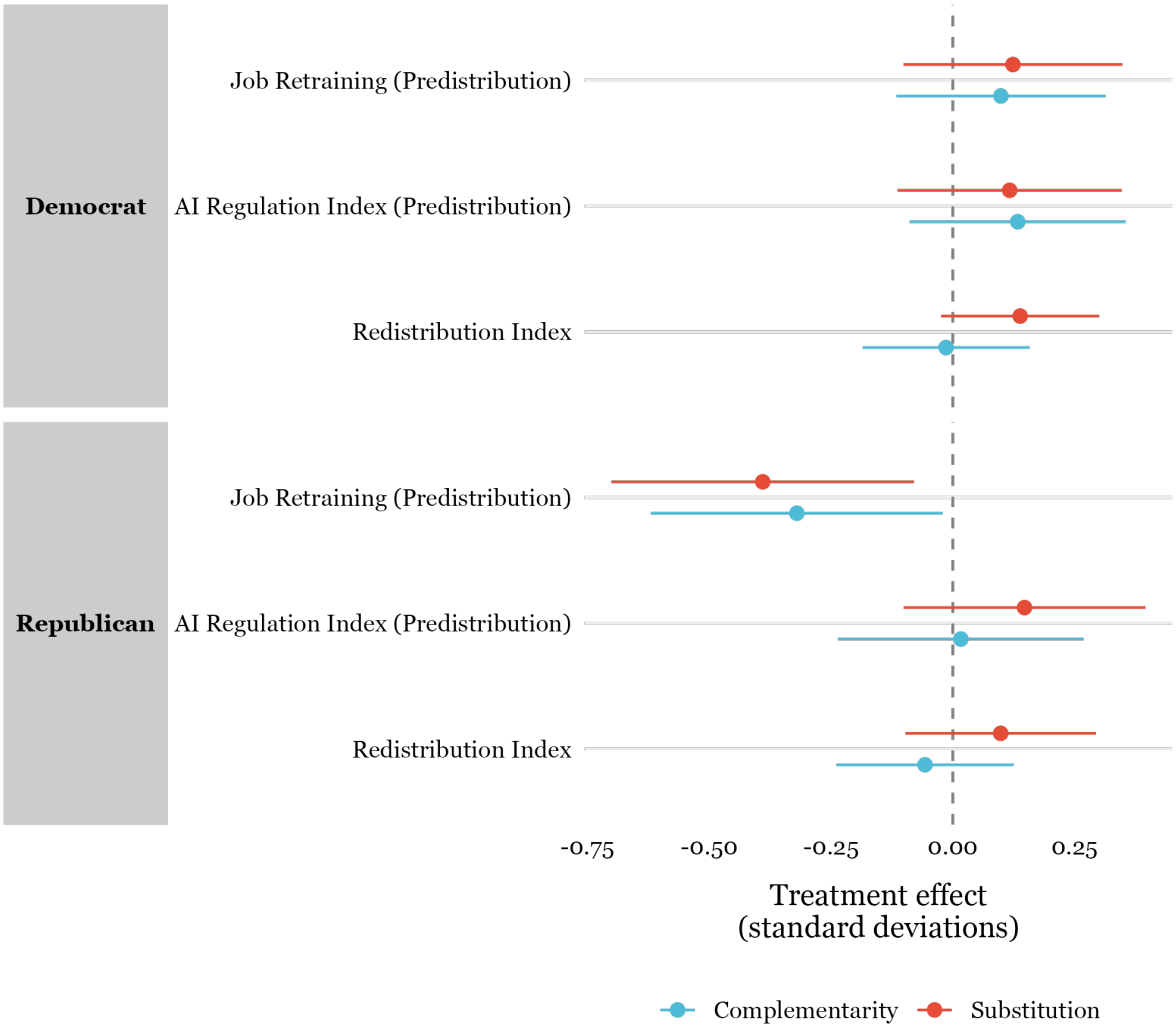


Figure 3: Treatment Effects by Party

The fact that Republicans in my sample respond to information about AI differently than Democrats begs the question of whether Republicans in this sample react to the AI threat differently in terms of their personal and sociotropic economic prospects. Republicans are especially responsive to AI-threat information, which dash their baseline optimism about the effects of AI on their incomes. Republicans in the control group expect AI to increase their incomes by 14.2%, but that is brought down to 1.4% by the substitution treatment. Additionally, Republicans who viewed the substitution treatment expect AI to reduce the total number of jobs. In tandem with downgrading both sociotropic and personal economic expectations, Republicans reduce support for predistribution in the form of job retraining programs for those laid off due to AI. For more details, see Appendix section C.2.

To address the concern that Republicans in my sample meaningfully differed from Democrats, I conducted a number of robustness checks. I begin with balance tests and confirmed that Democrats and Republicans in my sample are almost identical in sample composition on all demographics (age, gender, education level, ethnicity, AI risk). The largest difference between Democrats and Republicans in my sample is in income, and even that is quite small. Republicans in the sample earn an average of $4500 per month while Democrats earn an average of $4100 per month, a difference I argue is analytically negligible. Next, I ran a number of interaction and bootstrap models. Republican identity is a statistically significant moderator of reduced support for job retraining in response to both AI information treatments in both simple bootstrap and OLS interaction models, but not in fully saturated models (see Appendix Tables D.1 - D.6).

What exactly drives Republicans' changes in policy preferences? One explanation may be that Republicans simply work in differentially risky jobs with regard to AI. To test whether this is the case, I observe both the occupational risk score and potential of AI augmentation score from Loaiza and Rigobon (2024) and find Republicans and Democrats are nearly identical on both counts.[14] Another explanation is that Republicans and Democrats differ

[14]EPOCH Metrics of Average Partisan AI Risk and Augmentation potential in my sample: Republican AI Risk = .54 out of 1, Democrat AI Risk = .55; Republican AI Aug = .4, Democrat AI Aug = .4

in their *reactions* to occupational risk. To this argument, I offer suggestive evidence via an analysis using the Interflex package in R. Figures are available in Section D of the Online Appendix.

The Interflex analysis shows that personal occupational risk has a different relationship with support for redistribution and predistribution depending on one's partisan identity. For Republicans within the AI-threat group, marginal support for redistribution and predistribution decreases as AI-risk goes up, while the exact opposite is the case for Democrats. However the only statistically significant interaction of these Interflex analyses is that of between Republicans and redistribution support. Such a dynamic is not observed with regard to job retraining support, further adding doubt to whether it is differences in relationships with personal AI-risk between Republicans and Democrats that drive differences in AI-related policy preferences.

Ultimately, the goal of this study is not to identify exactly *why* Republicans and Democrats react differently to AI-threats, but rather, to show that they do. If political campaigns recognize that partisanship does indeed moderate perceptions of the AI threat as this study shows, we may be approaching the end of the era in which AI was a nonpartisan issue. (Gallego and Kurer, 2022)

Education Levels, the AI Threat, and Preferences for Predistribution

Given the amount of research highlighting the importance of education in predicting both vulnerability to automation and economic policy preferences for predistribution (e.g., (Autor et al., 2003; Hacker, 2011; Kuziemko et al., 2023; Vilbig, 2026) I test *H3*, which posits that the less educated prefer predistribution in response to negative AI frames. I find that those without a college degree do indeed respond to the AI threat by asking for increased regulation, a key form of predistribution. Those with degrees, on the other hand, do not exhibit any policy preference shifts. To some extent, these results are consistent with those of Haslberger et al. (2025), who show treated users with AI-salience by having them use ChatGPT in a

controlled setting. Post treatment, they show that those with a degree want less regulation in response to using AI. In other words, they show what might be considered an inverse effect of what I find, where degree holders are made to be less supportive of regulation post-treatment. As my sample is heavily weighted toward degree holders, I conduct a number of robustness checks on these differential results that are available in the Online Appendix, section E.

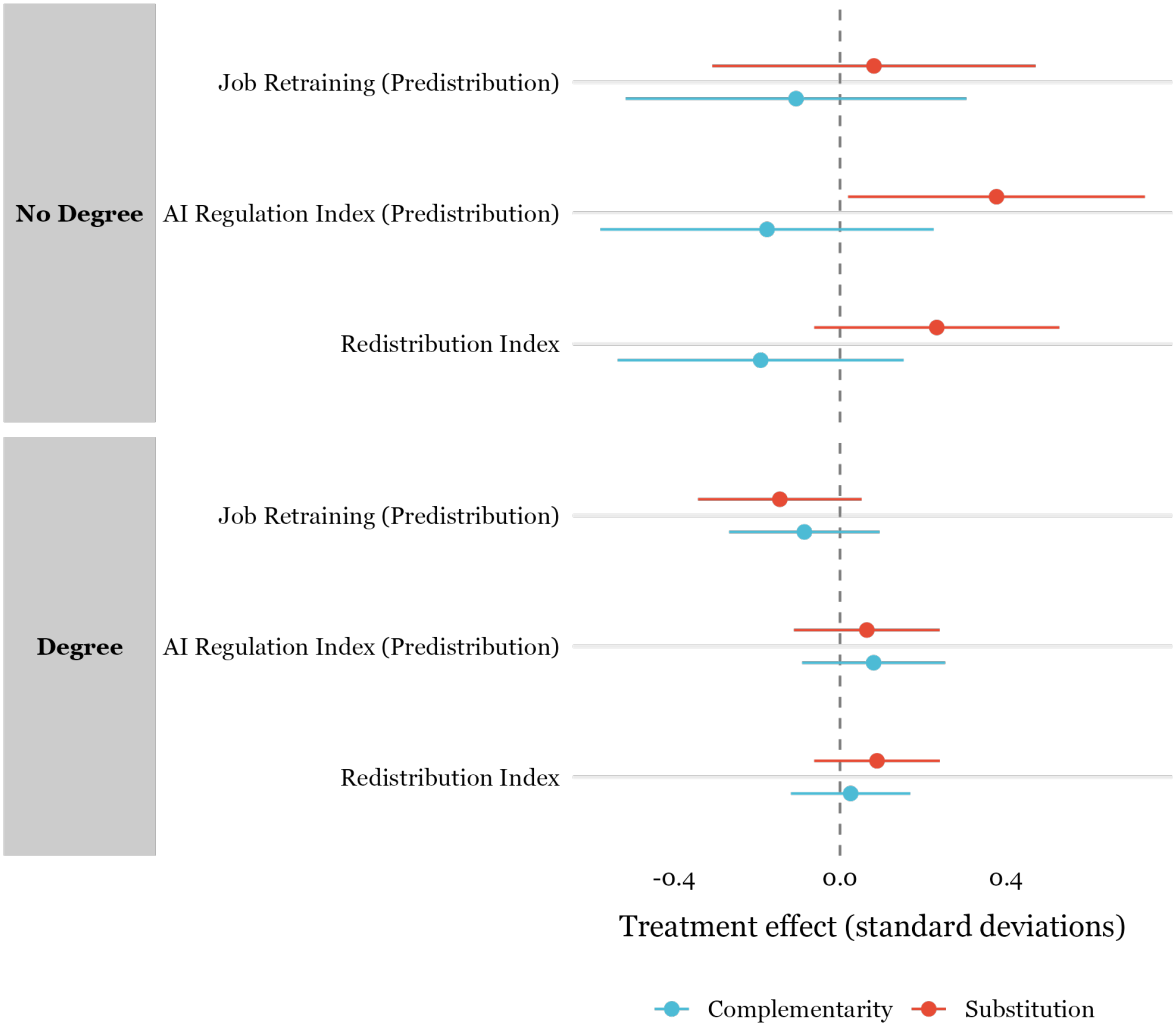


Figure 4: Results by Educational Attainment (AA included as a degree)

I next consider whether the less-educated differ in their predictions of AI's personal and sociotropic impacts post-treatment. I find that degree holders downgrade *both* their

personal income expectations and their projections of the total number of jobs in the economy in response to AI-threat information, but that non-degree holders only downgrade their sociotropic expectations. This result is consistent whether AA holders are considered degree holders or not. In fact, non degree holders are responsive to the complementarity treatment and assess themselves to be less likely to be unemployed by AI post-treatment. While individuals broadly expect AI to reduce the total number of jobs, these indicate suggest that non-degree holders are more likely to consider themselves immune to the broader labor market impacts of AI.

It is possible that education level is simply serving as a proxy for a respondents' AI vulnerability, which better explains these differences. I show this not to be the case, however. I find there to be a modest decrease in AI risk as educational attainment increases, but that the overall sample of white collar workers is nearly identical in terms of AI risk across education levels. A correlation test, though, reveals that ordinal increases in educational attainment are associated with a small decrease in AI risk (R = -.13) among these white-collar workers. This finding is somewhat surprising given the amount of research suggesting that degree holders are at greater risk of job displacement from AI because they are more likely to do jobs that can be completed on a computer (Felten et al., 2021). For example, a BA holding office worker is at more risk of AI-Chatbot-replacement than a GED holding electrician.

Table 1: Mean AI Risk (0-1 EPOCH Metric) by Education Level

Education Level	Mean
Some high school or less	0.56
High school diploma or GED	0.58
Some college, but no degree	0.57
Associates or technical degree	0.59
Bachelor's degree	0.55
Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, DDS etc.)	0.52

Some of this correlation between education level and AI risk is likely driven by the fact that that same BA holding office worker is more likely to be unemployed by AI than a graduate degree holding research scientist or doctor, as indicated by Table 3. More broadly, Table 3 shows that AI risk among white-collar workers is clustered together, suggesting that by virtue of being a white-collar worker, regardless of education level, one must live with a real level of AI risk. The most relevant implication of this is that despite similar levels of objective AI risk, the policy preferences of those without a degree are far more responsive to AI-threat frames than those with a degree.

Conclusion

A 2026 NBC poll shows that the vast majority of Americans have a negative opinion about AI, and that neither Democrats nor Republicans are satisfied with their party's policy response to AI (Smith, 2026). AI's labor market impact is an emerging political issue, but it is yet to be harnessed by a major party in the United States. It is in this context that white-collar voters find themselves awash with information about how AI will transform the white-collar economy for both good and bad, but previous research largely finds that this information is unable to convince voters to alter their economic policy preferences. One rea-

son for this, as the present study shows, is that voters are largely unconvinced by optimistic information about AI. In other words, white-collar voters have a persistent AI pessimism. The party which offers popular policy solutions to these workers may add an important bloc of voters to their coalition in the same way that displaced manufacturing workers in key swing states proved an important electoral gain for the Republican party in 2016 (Golin and Rauh, 2022; Broz et al., 2021; Colantone and Stanig, 2018; Magistro et al., 2025; Margalit, 2011). After all, public attitudes toward AI have similar microfoundations as those toward globalization, with significant tradeoffs between productivity increases and job losses concentrated in a particular sector (Magistro et al., 2025).

I show that voters do indeed prefer the government make policy changes in response to information about the economic impacts of AI, but that the specific policies they prefer vary widely based on key personal characteristics. As this study indicates, voters' policy demands are largely informed by party affiliation and education level.[15] I adopt the framework of Hacker (2011) and Kuziemko et al. (2023) in dividing preferred policy responses to "predistribution" and "redistribution", showing that it better explains the heterogeneous patterns we observe in AI-related policy preferences than models that view workers as either being pro or anti redistribution.

Many would expect increased redistribution to be the primary policy lever desired by voters facing information about the AI-threat, especially for those without a degree for whom a lost job may be more difficult to recover from. Yet, the present video-based survey experiment shows that those without a college degree prefer increased predistribution in the form of AI regulation in response to AI-threat information while those with degrees show no changes in policy preferences. This finding supports the predistribution vs. redistribution framework model in the AI-era, as it is consistent with the pioneering work of Kuziemko et al. (2023) and Hacker (2011) that argued predistribution was the key policy demand of

[15]Though not the focus of this study, a number of previous investigations show men and women react differently to automation threats. I conduct a subgroup analysis of this dynamic in the Online Appendix. In keeping with prior research, I show women are more supportive of regulating AI in the face of AI-threat primes.

those without college degrees in the pre-AI era. Given that those without college degrees are a voting bloc who have proven to be willing to realign in American presidential elections, this study shows that Democrats may be able to recapture some of their dwindling share of votes from this group by promising predistribution in the form of AI regulation.

Beyond this contribution, the present study offers an updated snapshot of the rapidly shifting partisan dynamic of AI in the US. While a conjoint experiment run by Magistro et al. (2025) shows that partisanship shapes support for AI-related public policy, it was conducted in 2023, before the formation of the Trump Administration's uneasy alliance with AI leaders. It is quite plausible that the public, friendly relationship between Silicon Valley AI-entrepreneurs like Peter Thiel, Elon Musk, and Alex Karp has changed the dynamic that Magistro et al. (2025) observe. With a video-based survey experiment conducted in late 2025 more than 7 months into Trump's second term, I show that Republicans respond to the labor market disruptions of AI differently than Democrats, preferring to reduce social safety nets when faced with both AI-threat and AI-complementarity frames. That these results differ from those of Magistro et al. (2025) suggests a potential a partisan shift in the perception of AI in the US since 2023.

Voters' information environments are saturated with descriptions of both the promise and peril of AI for the white-collar economy. If AI develops as promised, it may represent the most significant episode of creative destruction in modern history and carry with it tremendous political consequences. Yet, previous studies largely fail to show that information in this vein shifts voters' policy preferences despite research-based reasons to think that it should. Using a video-based survey experiment, this study shows that voters' preferred policy responses to AI are multidimensional and strongly informed by their party identification and education level. As Democrats gear up for the 2028 presidential election, this study suggests that offering AI regulation instead of redistribution may be an effective strategy in courting voters without a college degree.

References

Lene Aarøe and Michael Bang Petersen. Crowding Out Culture: Scandinavians and Americans Agree on Social Welfare in the Face of Deservingness Cues. *The Journal of Politics*, 76(3):684–697, July 2014. ISSN 0022-3816. doi: 10.1017/S002238161400019X. URL https://www.journals.uchicago.edu/doi/10.1017/S002238161400019X.

Daron Acemoglu and Simon Johnson. *Power and Progress: our thousand-year struggle over technology and prosperity*. 2023.

Daron Acemoglu and Pascual Restrepo. Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244, June 2020. ISSN 0022-3808. doi: 10.1086/705716. URL https://www.journals.uchicago.edu/doi/10.1086/705716.

David H. Autor, Frank Levy, and Richard J. Murnane. The Skill Content of Recent Technological Change: An Empirical Exploration*. *The Quarterly Journal of Economics*, 118 (4):1279–1333, November 2003. ISSN 0033-5533. doi: 10.1162/003355303322552801. URL https://doi.org/10.1162/003355303322552801.

Larry M Bartels. What's the Matter with What's the Matter with Kansas? 2006.

Roland Benabou and Efe A. Ok. Social Mobility and the Demand for Redistribution: The Poum Hypothesis*. *The Quarterly Journal of Economics*, 116(2):447–487, May 2001. ISSN 0033-5533. doi: 10.1162/00335530151144078. URL https://doi.org/10.1162/00335530151144078.

Carola Binder, Pascal Frank, and Jane M. Ryngaert. The Causal Effect of News on Inflation Expectations, August 2025. URL https://www.nber.org/papers/w34088.

Antoine Bozio, Bertrand Garbinti, Jonathan Goupille-Lebret, Malka Guillot, and Thomas Piketty. Predistribution versus Redistribution: Evidence from France and the United States. *American Economic Journal: Applied Economics*, 16(2):31–65, April 2024. ISSN

1945-7782. doi: 10.1257/app.20220023. URL https://www.aeaweb.org/articles?id=10.1257/app.20220023.

J. Lawrence Broz, Jeffry Frieden, and Stephen Weymouth. Populism in Place: The Economic Geography of the Globalization Backlash. *International Organization*, 75(2):464–494, February 2021. ISSN 0020-8183, 1531-5088. doi: 10.1017/S0020818320000314. URL https://www.cambridge.org/core/journals/international-organization/article/populism-in-place-the-economic-geography-of-the-globalization-backlash/98ED873D925E0590CB9A78AEC68BB439.

Marius R. Busemeyer and Alexander H. J. Sahm. Social Investment, Redistribution or Basic Income? Exploring the Association Between Automation Risk and Welfare State Attitudes in Europe. *Journal of Social Policy*, 51(4):751–770, October 2022. ISSN 0047-2794, 1469-7823. doi: 10.1017/S0047279421000519.

Roland Bénabou and Jean Tirole. Belief in a Just World and Redistributive Politics*. *The Quarterly Journal of Economics*, 121(2):699–746, May 2006. ISSN 0033-5533. doi: 10.1162/qjec.2006.121.2.699. URL https://doi.org/10.1162/qjec.2006.121.2.699.

Roland Bénabou and Jean Tirole. Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, 30(3):141–164, August 2016. ISSN 0895-3309. doi: 10.1257/jep.30.3.141. URL https://pubs.aeaweb.org/doi/10.1257/jep.30.3.141.

Daniel Checchi and Antonio Filippin. An Experimental Study of the Poum Hypothesis by Daniele Checchi, Antonio Filippin :: SSRN, 2003. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=464642.

CNBC. IBM CEO Arvind Krishna says A.I. will make more jobs than it will replace, May 2023.

Italo Colantone and Piero Stanig. Global Competition and Brexit. *American Political Science Review*, 112(2):201–218, May 2018. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055417000685. URL https://www.cambridge.org/core/journals/american-political-science-review/article/global-competition-and-brexit/C843990101DB9232B654E77130F88398.

Thomas R. Cusack, Torben Iversen, and Philipp Rehm. Economic Shocks, Inequality, and Popular Support for Redistribution. In Pablo Beramendi and Christopher J. Anderson, editors, *Democracy, Inequality, and Representation in Comparative Perspective*, pages 203–231. Russell Sage Foundation, 2008. ISBN 978-0-87154-088-1. URL http://www.jstor.org/stable/10.7758/9781610440448.11.

Rafaela M. Dancygier and Michael J. Donnelly. Sectoral Economies, Economic Contexts, and Attitudes toward Immigration. *The Journal of Politics*, 75(1), January 2013. ISSN 0022-3816. doi: 10.1017/S0022381612000849.

Thomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner. Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences. *Journal of the European Economic Association*, 9(3):522–550, June 2011. ISSN 1542-4766. doi: 10.1111/j.1542-4774.2011.01015.x. URL https://doi.org/10.1111/j.1542-4774.2011.01015.x.

James N. Druckman, Jordan Fein, and Thomas J. Leeper. A Source of Bias in Public Opinion Stability. *American Political Science Review*, 106(2):430–454, May 2012. ISSN 1537-5943, 0003-0554. doi: 10.1017/S0003055412000123. URL https://www.cambridge.org/core/journals/american-political-science-review/article/abs/source-of-bias-in-public-opinion-stability/D62AC6449C5FB464811427CC0C95F9AA.

Paul D. Ellis. *The Essential Guide to Effect Sizes: Statistical Power, Meta-*

Analysis, and the Interpretation of Research Results. Cambridge University Press, Cambridge, 2010. ISBN 978-0-521-19423-5. doi: 10.1017/CBO9780511761676. URL https://www.cambridge.org/core/books/essential-guide-to-effect-sizes/72C26CA99366A19CAC4EF5B16AE3297F.

Edward Felten, Manav Raj, and Robert Seamans. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217, 2021. ISSN 1097-0266. doi: 10.1002/smj.3286. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/smj.3286. _eprint: https://sms.onlinelibrary.wiley.com/doi/pdf/10.1002/smj.3286.

Christina Fong. Social preferences, self-interest, and the demand for redistribution. *Journal of Public Economics*, 82(2):225–246, November 2001. ISSN 0047-2727. doi: 10.1016/S0047-2727(00)00141-9. URL https://www.sciencedirect.com/science/article/pii/S0047272700001419.

Aina Gallego and Thomas Kurer. Automation, Digitalization, and Artificial Intelligence in the Workplace: Implications for Political Behavior. *Annual Review of Political Science*, 25(Volume 25, 2022):463–484, May 2022. ISSN 1094-2939, 1545-1577. doi: 10.1146/annurev-polisci-051120-104535. URL https://www.annualreviews.org/content/journals/10.1146/annurev-polisci-051120-104535.

Marcel Garz. News about the economy: a literature survey and methodological guidelines. *Oxford Open Economics*, 3:odae040, January 2024. ISSN 2752-5074. doi: 10.1093/ooec/odae040. URL https://doi.org/10.1093/ooec/odae040.

Andrew Gelman. Rich State, Poor State, Red State, Blue State: What's the Matter with Connecticut? *Quarterly Journal of Political Science*, 2(4):345–367, November 2007. ISSN 15540634. doi: 10.1561/100.00006026. URL http://www.nowpublishers.com/article/Details/QJPS-6026.

Spencer Goidel, Thiago MQ Moreira, and Brenna Armstrong. Party Realignment, Education, and the Turnout Advantage: Revisiting the Partisan Effect of Turnout. *American Politics Research*, 52(1):23–29, January 2024. ISSN 1532-673X. doi: 10.1177/1532673X231206149. URL https://doi.org/10.1177/1532673X231206149.

M. Golin and C. Rauh. The Impact of Fear of Automation. *Janeway Institute Working Papers*, December 2022. URL https://ideas.repec.org//p/cam/camjip/2229.html. Number: 2229.

Jane Green, Zack Grant, Geoffrey Evans, and Gaetano Inglese. Linking artificial intelligence job exposure to expectations: Understanding AI losers, winners, and their political preferences. *Research & Politics*, 12(2):20531680251337897, April 2025. ISSN 2053-1680. doi: 10.1177/20531680251337897. URL https://doi.org/10.1177/20531680251337897.

Scott Keeter and Ted Van Hannah Hartig Green, Andrew Daniller. Demographic profiles of Republican and Democratic voters, July 2023. URL https://www.pewresearch.org/politics/2023/07/12/demographic-profiles-of-republican-and-democratic-voters/.

Jacob Hacker. The institutional foundations of middle-class democracy. *Policy Network*, 6(5):33–37, 2011. URL https://scholar.google.com/scholar?cluster=11851519050439665 93&hl=en&oi=scholarr.

Matthias Haslberger, Jane Gingrich, and Jasmine Bhatia. Rage against the machine? Generative AI exposure, subjective risk, and policy preferences. *Journal of European Public Policy*, 0(0):1–28, September 2025. ISSN 1350-1763. doi: 10.1080/13501763.2025.2554903. URL https://doi.org/10.1080/13501763.2025.2554903. _eprint: https://doi.org/10.1080/13501763.2025.2554903.

Tobias Heinrich and Christopher Witko. Self-interest and preferences for the regulation of artificial intelligence. *Journal of Information Technology & Politics*, 22(3):306–321, July

2025. ISSN 1933-1681. doi: 10.1080/19331681.2024.2370815. URL https://doi.org/10.1080/19331681.2024.2370815. _eprint: https://doi.org/10.1080/19331681.2024.2370815.

Shanto Iyengar and Masha Krupenkin. Partisanship as Social Identity; Implications for the Study of Party Polarization. *The Forum*, 16(1):23–45, April 2018. ISSN 1540-8884, 2194-6183. doi: 10.1515/for-2018-0003. URL https://www.degruyter.com/document/doi/10.1515/for-2018-0003/html.

Sal Khan. Opinion | Sal Khan: A.I. Will Displace Workers at a Scale Many Don't Realize - The New York Times, 2025. URL https://www.nytimes.com/2025/12/27/opinion/artificial-intelligence-jobs-worker-training.html.

Ilyana Kuziemko, Nicolas Longuet-Marx, and Suresh Naidu. "Compensate the Losers?" Economic Policy and the Origins of U.S. Partisan Realignment, October 2023. URL https://www.nber.org/papers/w31794.

Melvin J. Lerner. The Belief in a Just World. In Melvin J. Lerner, editor, *The Belief in a Just World: A Fundamental Delusion*, pages 9–30. Springer US, Boston, MA, 1980. ISBN 978-1-4899-0448-5. doi: 10.1007/978-1-4899-0448-5_2. URL https://doi.org/10.1007/978-1-4899-0448-5_2.

Isabella Loaiza and Roberto Rigobon. The EPOCH of AI: Human-Machine Complementarities at Work, November 2024. URL https://papers.ssrn.com/abstract=5028371.

Beatrice Magistro, Sophie Borwein, R. Michael Alvarez, Bart Bonikowski, and Peter John Loewen. Attitudes toward artificial intelligence (AI) and globalization: Common microfoundations and political implications. *American Journal of Political Science*, 70(1):348–365, 2025. ISSN 1540-5907. doi: 10.1111/ajps.12959. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ajps.12959. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12959.

Yotam Margalit. Costly Jobs: Trade-related Layoffs, Government Compensation, and Voting in U.S. Elections. *American Political Science Review*, 105(1):166–188, February 2011. ISSN 1537-5943, 0003-0554. doi: 10.1017/S000305541000050X.

Yotam Margalit. Economic Insecurity and the Causes of Populism, Reconsidered. *Journal of Economic Perspectives*, 33(4):152–170, November 2019. ISSN 0895-3309. doi: 10.1257/jep.33.4.152. URL https://www.aeaweb.org/articles?id=10.1257/jep.33.4.152.

Yotam Margalit and Shir Raviv. The Politics of Using AI in Policy Implementation: Evidence from a Field Experiment, September 2023. URL https://papers.ssrn.com/abstract=4573250.

Anil Menon and Baobao Zhang. Future Shock or Future Shrug? Public Responses to Varied Artificial Intelligence Development Timelines. *The Journal of Politics*, October 2025. ISSN 0022-3816. doi: 10.1086/739200. URL https://www.journals.uchicago.edu/doi/10.1086/739200.

Dan Milmo. US jobs market yet to be seriously disrupted by AI, finds Yale study, 2025. URL https://www.theguardian.com/technology/2025/oct/01/us-jobs-market-yet-to-be-seriously-disrupted-by-ai-yale-study-chatgpt.

Daniela Morescalchi. Intel and Trump Administration Reach Historic Agreement to Accelerate American Technology and Manufacturing Leadership, August 2025. URL https://newsroom.intel.com/corporate/intel-and-trump-administration-reach-historic-agreement.

Diana C. Mutz. Improving Experimental Treatments in Political Science. In Donald P. Green and James N. Druckman, editors, *Advances in Experimental Political Science*, pages 219–238. Cambridge University Press, Cambridge, 2021. ISBN 978-1-108-47850-2. doi: 10.1017/9781108777919.016. URL https://www.cambridge.org/core/books/advances-in-experimental-political-science/

improving-experimental-treatments-in-political-science/9A348CD46C5D4D3DA9185BC764970680.

Nic Newman, Amy Ross Arguedas, Craig T. Robertson, Rasmus Kleis Nielsen, and Richard Fletcher. Reuters Insitute digital news report 2025. Technical report, Reuters Institute for the Study of Journalism, 2025. URL https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2025.

Kjell Noordzij, Willem de Koster, and Jeroen van der Waal. Explaining the educational gradient in trust in politicians: a video-vignette survey experiment. *West European Politics*, 47(4):784–812, June 2024. ISSN 0140-2382. doi: 10.1080/01402382.2023.2250163. URL https://doi.org/10.1080/01402382.2023.2250163. _eprint: https://doi.org/10.1080/01402382.2023.2250163.

Erica Owen and Noel P. Johnston. Occupation and the Political Economy of Trade: Job Routineness, Offshorability, and Protectionist Sentiment. *International Organization*, 71 (4):665–699, 2017. ISSN 0020-8183. URL https://www.jstor.org/stable/26569449.

Ryan Pendell. AI Use at Work Has Nearly Doubled in Two Years, June 2025. URL https://www.gallup.com/workplace/691643/work-nearly-doubled-two-years.aspx. Section: Workplace.

Martin Ravallion and Michael Lokshin. Who wants to redistribute?: The tunnel effect in 1990s Russia. *Journal of Public Economics*, 76(1):87–104, April 2000. ISSN 0047-2727. doi: 10.1016/S0047-2727(99)00064-X. URL https://www.sciencedirect.com/science/article/pii/S004727279900064X.

Claire Shefchik 3 min read. 'AI Can Write The Code': Zuckerberg Says Meta's Midlevel Engineers – Earning Six Figures – May Soon Be Replaced, January 2025. URL https://finance.yahoo.com/news/ai-write-code-zuckerberg-says-164519848.html.

Matt Schendstok and Sydney Schreiner Wertz. Occupational Exposure to Artificial Intelligence by Geography and Education. Technical report, US Department of Treasury, 2024.

MacKenzie Sigalos. MIT study finds AI can already replace 11.7% of U.S. workforce, November 2025. URL https://www.cnbc.com/2025/11/26/mit-study-finds-ai-can-already-replace-11point7percent-of-us-workforce.html. Section: Technology.

Allan Smith. Poll: Majority of voters say risks of AI outweigh the benefits, March 2026. URL https://www.nbcnews.com/politics/politics-news/poll-majority-voters-say-risks-ai-outweigh-benefits-rcna262196.

Alana Staiti. The Professional and Technical Workforce: By the Numbers, 2024. URL https://www.dpeaflcio.org/factsheets/the-professional-and-technical-workforce-by-the-numbers.

Charles S. Taber and Milton Lodge. Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science*, 50(3):755–769, 2006. ISSN 1540-5907. doi: 10.1111/j.1540-5907.2006.00214.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-5907.2006.00214.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-5907.2006.00214.x.

USDA. Trump Administration Announces $12 Billion Farmer Bridge Payments for American Farmers Impacted by Unfair Market Disruptions, 2025.

Stepan Vesely and Gloria Amaris. AI-driven income inequality and preferences for redistribution. *Economic Analysis and Policy*, 87:642–648, September 2025. ISSN 0313-5926. doi: 10.1016/j.eap.2025.06.019. URL https://www.sciencedirect.com/science/article/pii/S0313592625002462.

Karyn Vilbig. Why Do High-Income Democrats Support Redistribution? The Roles of Partisanship, Racial Attitudes and Fiscal Populism. *The British Jour-*

nal of Sociology, 77(1):30–41, 2026. ISSN 1468-4446. doi: 10.1111/1468-4446.70032. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-4446.70032. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-4446.70032.

Chloe Wittenberg, Ben M. Tappin, Adam J. Berinsky, and David G. Rand. The (minimal) persuasive advantage of political video over text. *Proceedings of the National Academy of Sciences*, 118(47):e2114388118, November 2021. doi: 10.1073/pnas.2114388118. URL https://www.pnas.org/doi/10.1073/pnas.2114388118.

Max Zahn. AI could create these new jobs despite gloomy forecasts, experts say, 2025. URL https://abcnews.com/Business/ai-create-new-jobs-despite-gloomy-forecasts-experts/story?id=123607557.

John R. Zaller. The Nature and Origins of Mass Opinion, August 1992. URL https://www.cambridge.org/core/books/nature-and-origins-of-mass-opinion/70B1485D3A9CFF55ADCCDD42FC7E926A. ISBN: 9780521404495 9780511818691 9780521407861.

Baobao Zhang. No Rage Against the Machines: Threat of Automation Does Not Change Policy Preferences. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 856–866, Oxford United Kingdom, July 2022. ACM. ISBN 978-1-4503-9247-1. doi: 10.1145/3514094.3534179. URL https://dl.acm.org/doi/10.1145/3514094.3534179.

Joshua N Zingher. TRENDS: Diploma divide: Educational attainment and the realignment of the American electorate. *Political Research Quarterly*, 75(2):263–277, June 2022. ISSN 1065-9129. doi: 10.1177/10659129221079862. URL https://doi.org/10.1177/10659129221079862.

Appendix: "Compensate the AI Losers?" A Video Experiment on AI-Framing and Policy Preferences

Contents

A Survey Text **1**

B Sample Information **24**
B.1 Demographic Approximations of U.S. White-Collar Workers age 25-54 24

C Subgroup Analyses Including Additional Policy Attitudes **25**
C.1 Full Sample 27
C.2 Partisanship 29
C.3 Education 36
C.4 Age 40
C.5 Sex 44
C.6 AI Knowledge 49
C.7 AI Self Efficacy 54
C.8 Social Mobility 59
C.9 AI Risk 66

D Robustness Checks: Results by Party **69**

E Robustness Checks: Results by Education Level **73**
E.1 Comparing Results by Education with Other Studies 82

F Average Causal Mediation Effects (ACME) **84**

G Other Analyses **89**
G.1 Bonferroni Correction P Values 89
G.2 Placebo ANOVA Test 89

A Survey Text

Title of Study: Economic and Technology Policy Preferences
Investigator: *Redacted for Manuscript Submission*

You are being invited to participate in a research study. If you agree to take part, you will be asked to provide demographic information, watch a short informational video, and complete a survey that asks your opinions on several political and economic issues. Your participation will take about 8–9 minutes. There is no direct benefit to you from taking part in this study. The risks are minimal. Some questions might make you feel uncomfortable or upset. You are free to skip any question or stop participating at any time.

Confidentiality
As with all research, there is a chance that confidentiality could be compromised; however, we are taking precautions to minimize this risk. Your responses will not include information that identifies you. Your Prolific ID will not be linked to your survey responses in any way that could identify you. It is used only to verify study completion and ensure you receive payment.

Compensation
You have been recruited to this study via Prolific and will receive the compensation stated in the Prolific posting upon completion of the study. To receive full payment, you must answer all questions and complete the study as described in the Prolific posting.

Voluntary Participation
Your participation in this study is voluntary. You may skip questions or stop participating at any time. However, full payment is contingent upon answering all questions and completing the study in full.

Questions
If you have questions about this research, you may contact the investigator at: [email address redacted]

Partial Disclosure
Research sometimes requires that information about its purpose not be fully disclosed to participants beforehand, because knowing the purpose could affect the results. None of the withheld information is reasonably expected to affect your willingness to take part. While the tasks you will be asked to perform have been described, the full intent of the research will be explained at the end of the study, at which point you may ask questions about the study's purpose and procedures.

Future Use of Data
We will use your information to conduct this study. Leftover data collected for this research may also be used for future research studies. These studies may be conducted by researchers at this institution or other institutions, including commercial entities. Data may be placed in one or more external scientific databases for access and use. We will not share any personally identifiable information, and we will not ask you for additional permission to share de-identified information. We may merge non-identifying prescreen information from Prolific (e.g.,

demographics you provided on Prolific) to your survey responses using a study ID. No personally identifying information will be stored with your responses.

Consent
By clicking the button below, you affirm that your participation in this survey is voluntary, that you are at least 18 years old, and that you understand you may skip questions or stop participating at any time, but full payment will only be provided if you answer all questions and complete the study in full.

- ○ I agree
- ○ I don't agree

Do you consider your current job to be "white-collar"? White-collar jobs are typically characterized by office work, administrative tasks, professional services, or managerial work, as opposed to primarily manual or physical labor.

- ○ Yes, my job is white-collar
- ○ No, my job is not white-collar
- ○ Not sure

What is your current or most recent occupation? Start typing and select the closest option from the dropdown menu.

Generally speaking, do you usually think of yourself as a Republican, a Democrat, an Independent?

- ○ Republican
- ○ Democrat
- ○ Independent
- ○ Other _____
- ○ No preference

What is the highest level of education you have completed?

- ○ Some high school or less
- ○ High school diploma or GED
- ○ Some college, but no degree
- ○ Associates or technical degree
- ○ Bachelor's degree
- ○ Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, DDS etc.)
- ○ Prefer not to say

What is your average monthly income? A rough estimate is acceptable.

_____

How would you describe your socioeconomic class growing up?

- ○ Upper class
- ○ Upper middle class
- ○ Middle class
- ○ Lower middle class
- ○ Lower class
- ○ I don't know

What socioeconomic class do you expect to be in the future?

- ○ Upper class
- ○ Upper middle class
- ○ Middle class
- ○ Lower middle class
- ○ Lower class
- ○ I don't know

When interacting with artificial intelligence (AI) technologies/products, I feel very relaxed.

- ○ Exactly true
- ○ Moderately true
- ○ Hardly true
- ○ Not at all true
- ○ Not sure

I can happily interact with AI technologies/products smoothly.

- ○ Exactly true
- ○ Moderately true
- ○ Hardly true
- ○ Not at all true
- ○ Not sure

AI technologies/products jargon does not baffle me.

- ○ Exactly true
- ○ Moderately true
- ○ Hardly true
- ○ Not at all true
- ○ Not sure

AI technologies/products are good aids to learning.

- ○ Exactly true
- ○ Moderately true
- ○ Hardly true
- ○ Not at all true
- ○ Not sure

AI technologies/products help me to save a lot of time.

- ○ Exactly true
- ○ Moderately true
- ○ Hardly true
- ○ Not at all true
- ○ Not sure

I find it easy to get AI technologies/products to do what I want it to do.

- ○ Exactly true
- ○ Moderately true
- ○ Hardly true
- ○ Not at all true
- ○ Not sure

How do you see yourself: are you generally a person who is fully prepared to take risks or do you try to avoid taking risks? Please tick a box on the scale, where the value 0 means: 'not at all willing to take risks' and the value 10 means: 'very willing to take risks'.

Not at all willing to take risks Very willing to take risks

0 1 2 3 4 5 6 7 8 9 10



This is an attention check. Please select "Moderately important"

- Not at all important
- Slightly important
- Moderately important
- Very important
- Extremely important

What is "OpenAI"?

- An organization lobbying for "open borders"
- A technology company
- An organization lobbying for open, non-discriminatory elections
- An organization lobbying for open access to government documents
- Other
- I don't know

Large language models like ChatGPT and Claude are primarily trained to do what?

- ○ Search the internet in real time
- ○ Perform mathematical proofs
- ○ Translate only between languages
- ○ Predict the next word in a sequence of text

What does "RLHF" stand for?

- ○ Reinforcement Learning from Human Feedback
- ○ Risky Lifestyle Health Factors
- ○ Road Less Traveled Hiking Federation
- ○ Renewable Liquid Hydrogen Fuel
- ○ Other
- ○ I don't know

Who is Sam Altman?

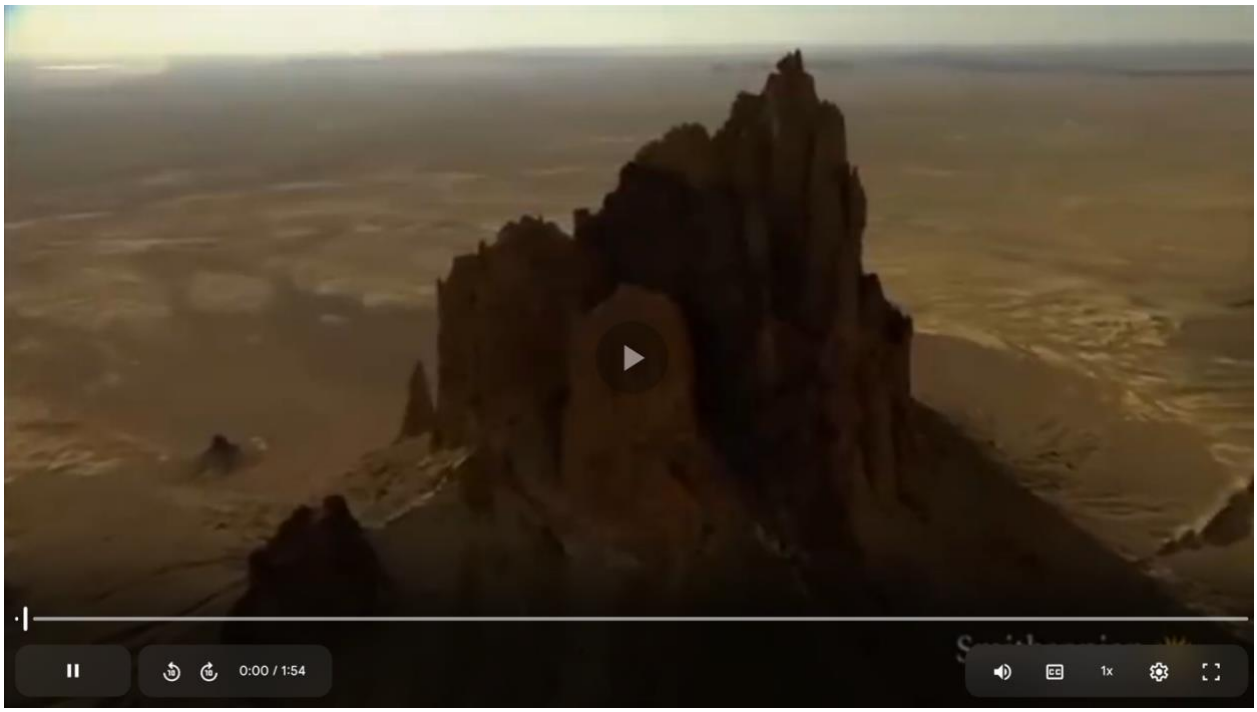
- ○ Member of President Trump's cabinet
- ○ NCAA football star
- ○ Influential technology entrepreneur
- ○ Academy Award-nominated screenwriter
- ○ Other
- ○ I don't know

At this point in the survey, respondents were randomly sorted into one of three groups:

1) **Control Treatments (1/3 of total respondents)**
    a. **¼ saw the Grand Canyon video**
    b. **¼ saw the Pyramids of Giza video**
    c. **¼ saw the Juggling video**
    d. **¼ saw the Giraffes video**
2) **Substitution Treatment Video (1/3 of total respondents)**
3) **Complementarity Treatment Video (1/3 of total respondents)**

<u>Each respondent only saw one of the following videos.</u>

Please watch this short informational video. Once you finish watching the video, you will be able to go to the next page.



The Grand Canyon (1:54)

<u>Video Transcript (not shown to respondents unless they turned on captions):</u> In the early 19th century, most European scientists believed that volcanoes and earthquakes were the primary forces that created the landscape around us, from the highest mountains to the deepest valleys and the vast open spaces that lie in between. But in 1858, an American geologist named John Newberry changed that thinking after he arrived in the

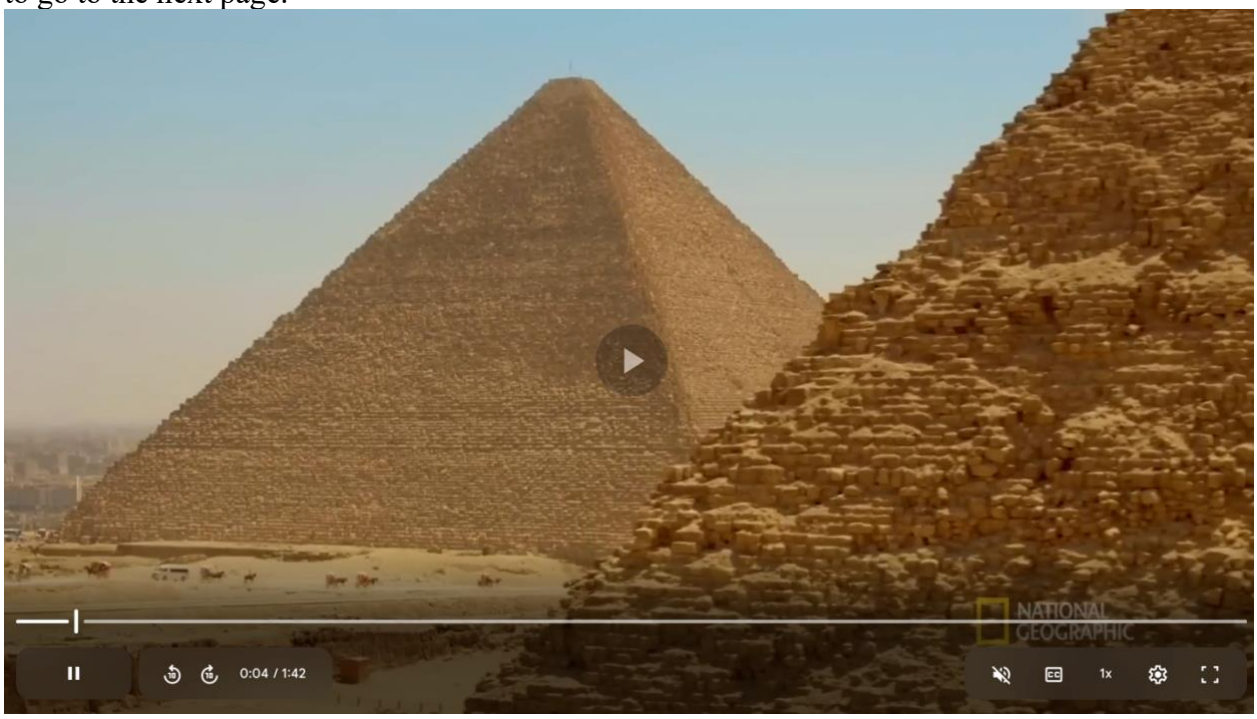
Grand Canyon. He came on a surveying expedition for the US government. At the time, the Grand Canyon didn't even appear on many maps. It was part of a region written off simply as the Great Unknown.

Newberry was the first scientist ever to descend down to the canyon's floor. Once he set eyes on the Colorado River and studied the colorful layers of rock that towered above him, he proposed that the river itself had carved the canyon. It was a radical theory for the middle of the 19th century. It suggested that it would have taken millions of years to carve the canyon, which would mean that the Earth had to be much older than most people imagined. Today, most scientists agree that Newberry was right about how the canyon was formed. It all began about 6 million years ago when the Colorado River first began flowing down and out of the Rocky Mountains and started carving a path through the Colorado Plateau.

What was the video you just watched about?

- ○ The Grand Canyon
- ○ The Pyramids of Giza
- ○ The eating habits of parrots
- ○ How to juggle

Please watch this short informational video. Once you finish watching the video, you will be able to go to the next page.



Video Transcript (not shown to respondents unless they turned on captions): ***Pyramids of Giza (1:37)***

The west bank of the river Nile is home to the world's most iconic monuments: the mighty pyramids of Giza. The pyramids once housed the bodies of the pharaohs, but though ancient Egyptian civilization lasted for nearly 3,000 years, its kings only built huge tombs like these for a few centuries. Egyptologists are still trying to piece together why the pharaohs stopped constructing giant pyramids.

For Egyptologist Chris Norton, the majesty of the ancient structures makes the fact that Egyptians gave up building them all the more incredible. Ten miles south of the legendary pyramids of Giza is Saqqara. When we think about pyramids, we tend to think of Giza and the Great Pyramid of Khufu in particular, but actually this is where it all began. Chris has come to the birthplace of pyramid building to search for clues to why Egyptians built giant pyramids for less than 500 years.

Constructed a century before the iconic pyramids at Giza, Egypt's first pyramid is a 200-foot-tall mausoleum of six huge limestone platforms, carefully engineered to spread the weight of rock and prevent collapse. Deep inside is a giant shaft 26 feet wide and 82 feet deep. At the bottom lies the intended final resting place of the pharaoh Djoser.

What was the video you just watched about?

- ○ The Grand Canyon
- ○ The Pyramids of Giza
- ○ The eating habits of parrots
- ○ How to juggle

Please watch this short informational video. Once you finish watching the video, you will be able to go to the next page.



Video Transcript (not shown to respondents unless they turned on captions): ***Giraffe Video (1:53)***

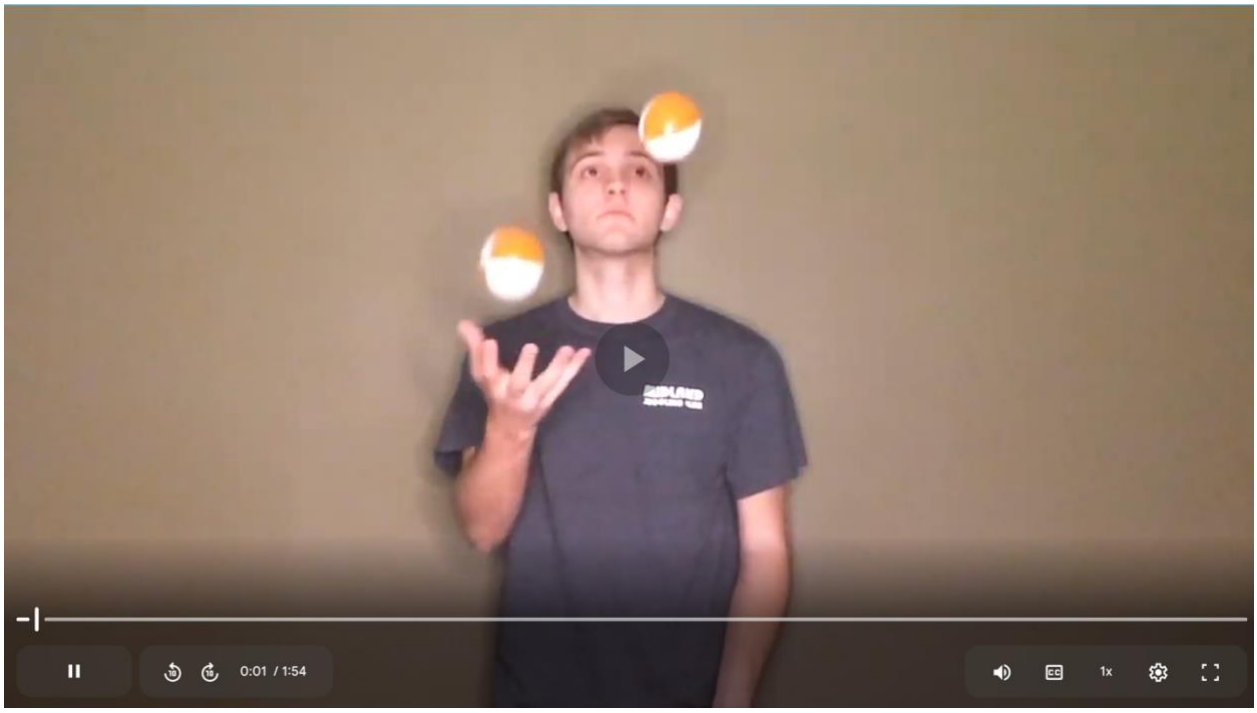
With their towering height and signature long necks, giraffes are some of the most iconic creatures of Africa. Giraffes are the tallest living land animals. Adult females can grow over 14 feet tall and adult males can reach over 18 feet. Baby giraffes, called calves, are born 6 feet tall and may grow an additional inch every day during their first week. This exceptional height allows giraffes to eat leaves and buds from treetops unreachable by other animals and to look out for predators and other oncoming dangers.

Giraffes have exceptionally strong hearts. Their hearts are about 2 feet long and weigh approximately 25 pounds, equal to about 50 human hearts. Because of the great distance between a giraffe's heart and its brain due to the animal's long neck, its heart has evolved thick muscular walls. Thanks to these powerful walls, a giraffe's blood pressure can be two times higher than many other animals. This helps the giraffe's heart defy gravity and push blood up its long neck and into its brain.

Giraffes have the same number of neckbones as humans. Both humans and giraffes only have seven bones in their necks. However, human neckbones are only half an inch tall while giraffe bones are 10 inches tall. This means giraffe necks can grow to be about 6 feet long and weigh 600 pounds.

What was the video you just watched about?

- ○ Giraffes
- ○ The Pyramids of Giza
- ○ The eating habits of parrots
- ○ How to juggle



Video Transcript (not shown to respondents unless they turned on captions): ***Juggling Video (1:54)***

Today I'm going to be showing you how to juggle three balls in just two minutes. It will take a lot longer than a few minutes to learn, but I will be outlining all the steps you need to know. First, start with just one ball. Put your feet shoulder-width apart, hands at your sides, and look slightly upwards. Do not watch the ball go into your hand but instead focus on the ball at the peak, and your hands will know where to go. This may seem trivial for many of you, but keep practicing until it feels solid and comfortable.

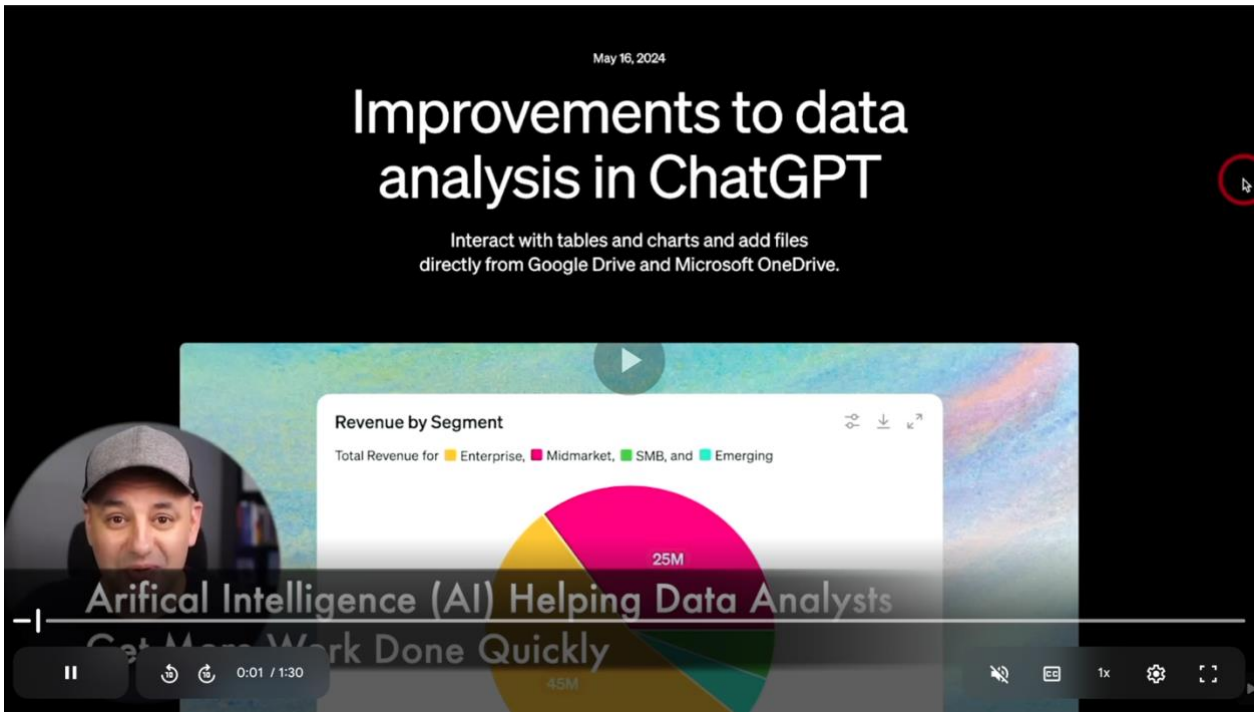
Now grab your second juggling ball. Many of you will want to do what's called a shower pattern, where one hand does all the throwing and the balls are passed from hand to hand. This is not what you want to do. Instead, throw both balls into the air starting with your dominant hand and releasing the second ball when the first ball is at its peak. The balls will sort of make an X pattern in the air. This may seem more difficult, but it is essential to learn to be able to juggle three balls. Practice this until you can do it consistently without dropping.

Starting from both hands, grab your third juggling ball. Start with two juggling balls in your dominant hand and one in the other. Your first throw will be made with your dominant hand. You will do exactly what you did before when you juggled two, except now as the second ball reaches its peak you throw the third ball. Try and keep a similar tempo to what you see in the video. Many will try and rush the throws when they first start off. Make sure to keep your feet still and make sure your throws do not go too far forward or backwards. If you have trouble with your throws, try and go back to practicing two for a little while and make sure your throws are very clean. Start adding more throws once you feel comfortable doing three throws and catches, and slowly work your way up until you can run the pattern for a longer time. And just like that, you're juggling. Hopefully you found this helpful. Thank you for watching.

What was the video you just watched about?

- ○ Giraffes
- ○ The Pyramids of Giza
- ○ The eating habits of parrots
- ○ Juggling

Several studies report that workers who adopt Artificial Intelligence (AI) tools can complete tasks faster and take on higher-value work. These analyses estimate **potential pay increases of 20–30% for white-collar workers, alongside growth in overall employment in AI-exposed occupations.** Please watch this short video showing how AI tools like ChatGPT **are helping white collar workers earn more money** by helping them complete their tasks faster. Once you finish watching the video, you will be able to go to the next page.



Complementarity Treatment (1:30)

<u>Video Transcript (not shown to respondents unless they turned on captions):</u> ChatGPT just rolled out one of the most useful updates I've ever seen, and it's about data analysis. Something that would take me two hours I did in four minutes. Basically, all I did was give it a document from my drive and ask for a PowerPoint presentation. I opened up the PowerPoint and just used the AI designer to make some alterations to make things look nice. It created all these different graphs and analyses for me, and this maybe took me three to four minutes. The data analysis and creating these charts was all done with ChatGPT.

It also gave me this: I asked it to pretend it was an attorney representing my client in a negotiation over the termination of services, and it created bullet points of the different things I wanted to convey in the response. The response was tone-perfect and done in a much shorter period of time.

For course and lesson plan creation, use Education Copilot, ChatGPT, or ClassPoint. No longer will we have to spend hours designing lesson material or creating visually stunning and engaging content for lessons. Use ChatGPT to ask for a lesson plan or a rubric, or use ClassPoint to generate AI-powered quiz questions right inside PowerPoint based on the information written on your slide.

What was the video you just watched about?

- ○ How AI helped discover a new species of bacteria
- ○ The Pyramids of Giza
- ○ How AI can help white collar workers be more productive and earn more money
- ○ The Grand Canyon

Several recent studies warn that companies may use Artificial Intelligence (AI) to automate a substantial share of professional work. **These analyses estimate that about 30% of white-collar workers could face wage reductions or job loss** if firms replace tasks previously done by people with AI. Please watch this short informational video showing how **AI** platforms like ChatGPT, Claude, and Gemini are **already causing layoffs and reducing wages for white collar workers** because of AI's ability to do work cheaply and quickly. Once you finish watching the video, you will be able to go to the next page.



Video Transcript (not shown to respondents unless they turned on captions): ***AI Substitution Treatment (1:49)*** *Here is how AI is already taking over tasks at white collar jobs. For some tech companies, AI is contributing to job cuts. A May report said nearly 4,000 tech workers lost their jobs to AI. Artificial intelligence is also eliminating future positions. Take IBM as an example. Its CEO said the company could replace workers with AI. "I do believe that AI is going to replace a lot of what I'm calling white collar clerical jobs, the ones that are much more repetitive, the ones that people do the same task again and again and again. I think a good 30 percent of those roles could go away over five years."*

AI's role in television and film production is a key point of negotiations between the Writers Guild of America and the studios. The co-showrunner for The Handmaid's Tale laid out the writer's union's concerns: "We could deduce that that plan was probably one writer in a room with one or two machines creating an entire show. If it's coming for us, the creators of imaginary worlds, it is literally going to come for everyone."

UPS is the latest company announcing they're cutting 12,000 jobs, eliminating more than two percent of its workforce around the world. The move is part of an effort to save money, including turning to artificial intelligence to replace some of those jobs. Companies such as UPS are the latest to make sweeping job cuts, eliminating more than two percent of its global workforce in an effort to save costs, including turning to AI to replace some of the jobs. UNC's Mark McNealy says that the types of jobs at risk include customer service, as well as the legal area. As for UPS, he expects that managers could make up most of the layoffs.

What was the video you just watched about?

- ○ How AI helped discover a new species of bacteria
- ○ The Pyramids of Giza
- ○ How AI is threatening job security and reducing wages for white-collar workers
- ○ The Grand Canyon

What is the chance you estimate you will be unemployed due to AI driven job automation in the next 5 years?

- ○ Very likely to be unemployed due to AI
- ○ Somewhat likely to be unemployed due to AI
- ○ AI will have no effect on my employment status
- ○ Somewhat unlikely to be unemployed due to AI
- ○ Very unlikely to be unemployed due to AI

Please estimate the percentage impact that artificial intelligence will have on your personal income. Use negative values for a decrease and positive values for an increase.

-100% (Large decrease in income) | 0% (No effect) | 100% (Large increase in income)

-100 -80 -60 -40 -20 0 20 40 60 80 100



Thinking about developments in artificial intelligence, what effect, if any, do you think AI will have on the total number of jobs in the economy overall?

- ○ It will greatly decrease the number of jobs
- ○ It will somewhat decrease the number of jobs
- ○ It will have no major effect on the number of jobs
- ○ It will somewhat increase the number of jobs
- ○ It will greatly increase the number of jobs

Some people feel that the government should increase taxes a lot and spend much more on health and social services such as unemployment benefits and low-cost healthcare programs like Medicaid. Other people feel that the government should cut taxes a lot and spend much less on

health and social services. And other people have views somewhere in-between. Please use the slider to indicate where your views fall on this scale.

Government should cut taxes a lot and spend much less on health and social services

Government should increase taxes a lot and spend much more on health and social services

0 10 20 30 40 50 60 70 80 90 100



What should happen to federal tax rates on household income over $400,000? Use the slider to indicate your preference. 0 = Lower rates a lot 50 = Keep about the same 100 = Raise rates a lot

Lower rates a lot Raise rates a lot

0 10 20 30 40 50 60 70 80 90 100



What should happen to federal tax rates on large tech corporations? Use the slider to indicate your preference. 0 = Lower rates a lot 50 = Keep about the same 100 = Raise rates a lot

Lower rates a lot Raise rates a lot

0 10 20 30 40 50 60 70 80 90 100



The government should fund job retraining programs for workers who lose their jobs due to automation or artificial intelligence.

- ○ Strongly agree
- ○ Somewhat agree
- ○ Neither agree nor disagree
- ○ Somewhat disagree
- ○ Strongly disagree

How interested are you in joining a labor union?

- ○ Strongly interested
- ○ Moderately interested
- ○ Neither interested nor uninterested
- ○ Slightly interested
- ○ Not at all interested
- ○ I am already a member of a labor union

There is a discussion about the introduction of universal basic income. Universal basic income is a concept where everybody, regardless of whether they work or not, receives the same amount of money every month from the government. If universal basic income were to be introduced, how much should everybody receive per month? Put 0 If you don't want UBI to be introduced.

Dollar amount per month (0 – 10,000)	

Do you think the government should stop companies from replacing workers with Artificial Intelligence?

- ○ Strongly agree
- ○ Somewhat agree
- ○ Neither agree nor disagree
- ○ Somewhat disagree
- ○ Strongly disagree

Do you believe that the federal government should require companies to implement safety measures for new technologies, even if they increase the cost of products or services?

- ○ Strongly agree
- ○ Agree
- ○ Neither agree nor disagree
- ○ Disagree
- ○ Strongly disagree

Some people propose creating a new federal agency that oversees and licenses major artificial intelligence projects. What would you think about this idea?

- ○ Strongly agree
- ○ Agree
- ○ Neither agree nor disagree
- ○ Disagree
- ○ Strongly disagree

Do you think the federal government should do more or less to regulate companies developing artificial intelligence products?

- ○ A lot more
- ○ More
- ○ Neither more nor less
- ○ Less
- ○ A lot less
- ○ Not sure

Thank you for taking part in this survey! Please click the button below to be redirected back to Prolific and register your submission.

B Sample Information

B.1 Demographic Approximations of U.S. White-Collar Workers age 25-54

As AI's potential to transform work is most salient for white-collar workers, the survey experiment was run exclusively on a panel of white-collar workers. To ensure the survey sample was representative of the U.S. white-collar workforce, approximated the demographic breakdown of white-collar workers from ages 25-54 using US Census data and then set quotas for that precise demographic breakdown with Prolific. The US Census provides the Equal Employment Opportunity tabulation (most recently for 2013-2018) on occupational groups by Sex and Ethnicity based on subfield of employment. I specifically measured the demographic breakdown of white-collar workers by selecting only for individuals with those occupations in my demographic estimation.

Below, I list which US Census designated occupations I consider to be white collar. I use quotas based on demographics of people 16 and up, since racial and gender breakdown of white collar workers ages 25 to 54 is not readily available. I follow the EEO classification scheme: Hispanic or Latino (any race); and among non-Hispanics: White alone, Black alone, AIAN alone, Asian alone, NH/PI alone, and a residual 'Non-Hispanic Multiracial/Other' ('balance') comprising non-Hispanic respondents of two or more races and 'some other race'. To approximate the partisan breakdown of the U.S. white-collar workforce, I used data from the American National Election Survey (ANES) and Bureau of Labor Statistics (BLS).

White Collar

- Science, engineering and computer professionals
- Healthcare practitioner professionals
- Other professional workers
- Technicians
- Sales workers
- Administrative support workers

Non White-Collar

- Construction and extractive craft workers
- Installation, maintenance and repair craft workers
- Production operative workers
- Transportation and material moving operative workers
- Laborers and helpers
- Protective service workers
- Service workers, except protective

Using Prolific's quota-setting feature, I recruited a sample of 1009 individuals that mirrors the U.S. white-collar workforce in terms of partisanship, age, sex, and ethnicity. A post-treatment attention check confirmed comprehension, and participants who failed either of two attention checks were excluded from analysis. Of the 1009 respondents, 12 failed the attention checks. To limit Prolific's recruiting to white-collar workers, I set an initial screening criteria that only allowed recruitment of respondents who spend 50% or more of their time at work on a computer. As an additional step to ensure my sample was fully white-collar, I included a questions asking if the respondent considered themselves white-collar and asked for their occupation, which was matched with SOC-18 occupation codes. 60 respondents claimed they were not white-collar workers and 15 said they were not sure. I independently checked occupation codes for all of 75 of these workers and found that 45 of them were not white-collar workers. Those 45 respondents were dropped from the analysis.

Table B.1: U.S. White-Collar Workforce Demographics (EEO-ALL03W, 2014–2018, Age 16+

Dimension	Group	Count	Percent
Race/Ethnicity	Hispanic or Latino (any race)	9,832,660	11.7%
Race/Ethnicity	White alone (Not Hispanic)	58,502,815	69.5%
Race/Ethnicity	Black or African American alone (Not Hispanic)	7,878,330	9.4%
Race/Ethnicity	Asian alone (Not Hispanic)	5,757,095	6.8%
Race/Ethnicity	American Indian/Alaska Native alone (Not Hispanic)	356,000	0.4%
Race/Ethnicity	Native Hawaiian/Other Pacific Islander alone (Not Hispanic)	113,030	0.1%
Race/Ethnicity	Balance of Not Hispanic (e.g., Two+ races)	1,683,635	2.0%
Sex	Male	37,529,320	44.6%
Sex	Female	46,594,245	55.4%

Table B.2: Prolific Sample Sex

Category	Count	%
Female	557	54.70%
Male	462	45.30%

Table B.3: Prolific Sample Political Affiliation

Category	Count	%
Democrat	423	41.50%
Independent	186	18.30%
Other	43	4.20%
Republican	367	36.00%

Table B.4: Employment Status

Category	Count	%
Full-Time	852	83.60%
Part-Time	167	16.40%

Table B.5: Prolific Sample Ethnicity Distribution

Category	Count	%
Black/African American	93	9.10%
East Asian	22	2.20%
Latino/Hispanic	121	11.90%
Mixed	10	1.00%
Native American or Alaskan Native	7	0.70%
Other	7	0.70%
South Asian	17	1.70%
South East Asian	18	1.80%
White	693	68.00%
NA	31	3.00%

C Subgroup Analyses Including Additional Policy Attitudes

In this section I show treatment effects on additional AI-related attitudes which include interest in joining a union, support for universal basic income, and the components of the redistribution and regulation indices.

Table C.1: Control Group Baseline Descriptive Statistics: Sample Predisposition Prior to Treatment

Outcome	**Mean**	**SD**	**N**
AI Regulation Index (Predistribution)	0.73	0.76	320
Redistribution Index	16.34	20.37	321
Job Retraining (Predistribution)	1.21	0.89	320
Personal AI-Unemployment Concern (Egoistic)	−0.38	1.26	320
Expected % Impact of AI on Income (Egoistic)	4.77	36.78	321
Expected Impact of AI on Total # of Jobs (Sociotropic)	−0.77	0.96	319

Note: All variables scaled -2 - 2 except for Redistribution Index, which is scaled -50 - 50

C.1 Full Sample

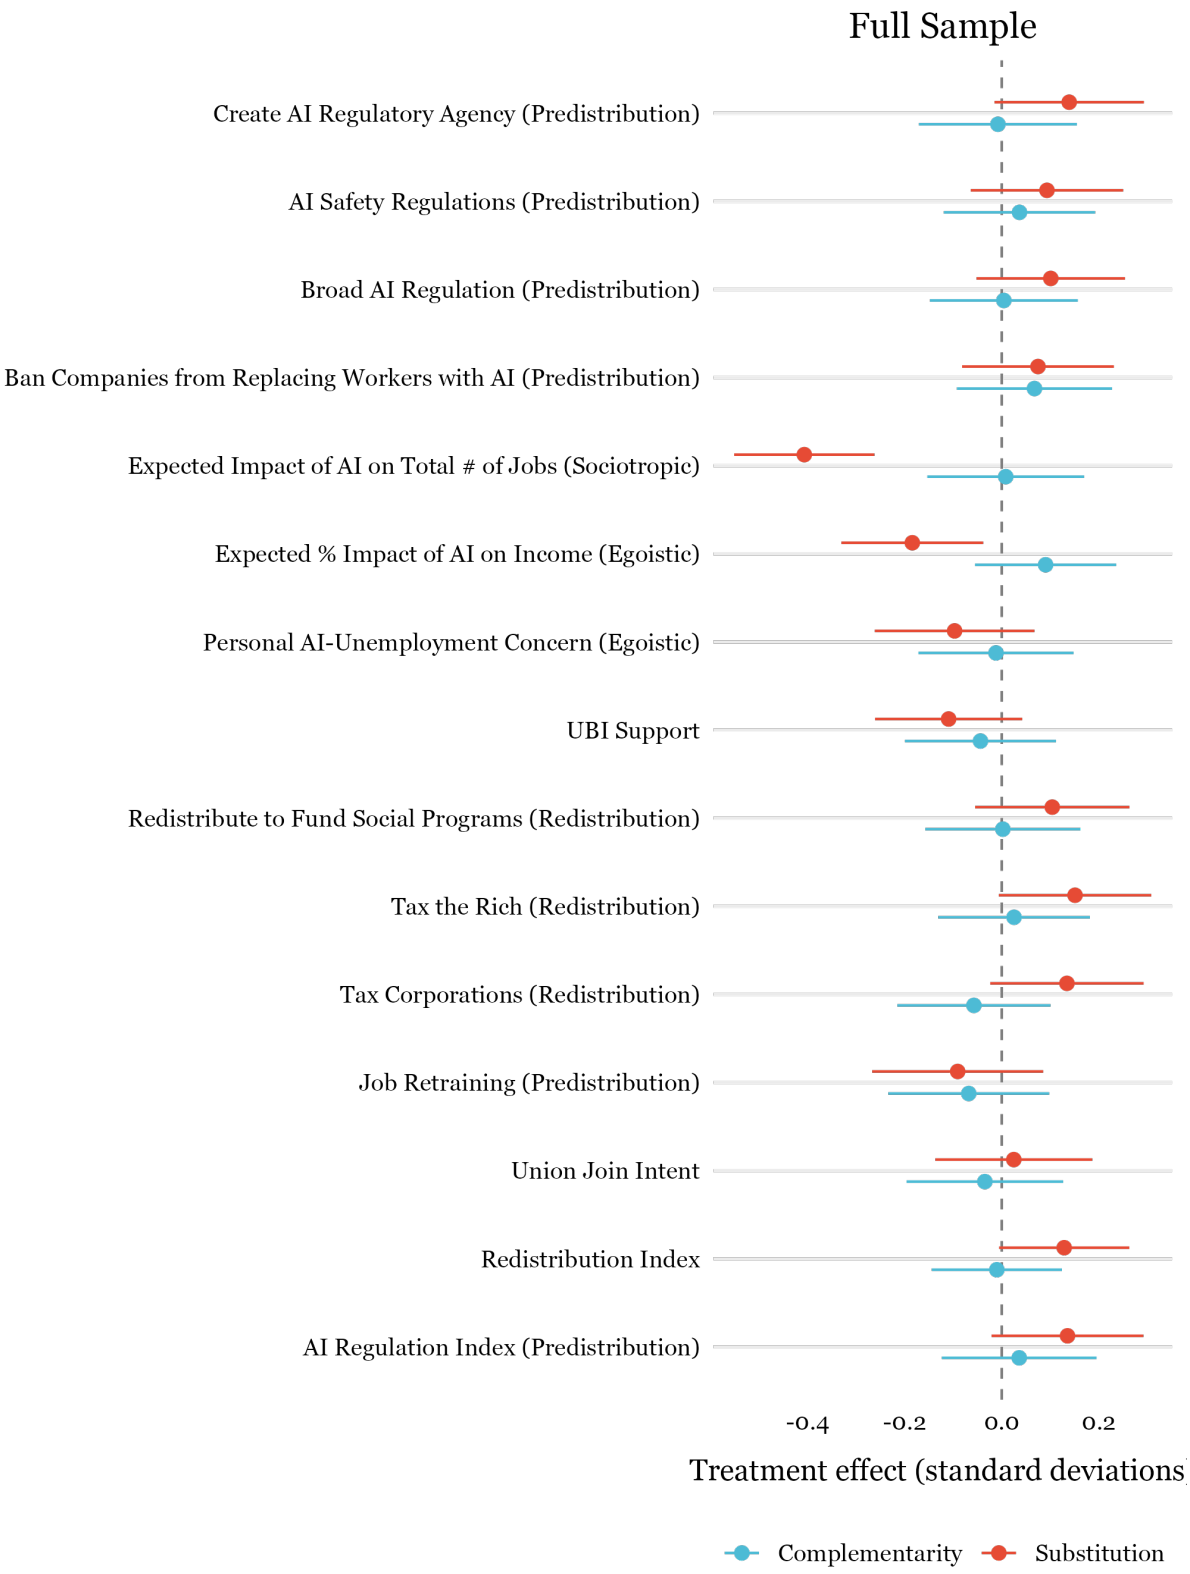


Table C.2: Full Sample: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.129+	−0.109	−0.091	0.135+	0.151+	0.134+	0.025
	(0.068)	(0.077)	(0.090)	(0.080)	(0.080)	(0.081)	(0.082)
Complementarity	−0.010	−0.044	−0.068	0.036	0.025	−0.057	−0.035
	(0.068)	(0.079)	(0.085)	(0.081)	(0.080)	(0.080)	(0.082)
Num.Obs.	929	917	927	927	928	928	898
RMSE	0.85	0.96	1.11	1.00	0.99	1.00	1.01

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.3: Full Sample: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.101	0.075	0.093	0.139+
	(0.078)	(0.080)	(0.080)	(0.078)
Complementarity	0.004	0.067	0.037	−0.008
	(0.078)	(0.081)	(0.079)	(0.083)
Num.Obs.	917	927	927	926
RMSE	0.95	1.00	0.99	1.00

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.4: Full Sample: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−6.777*	−0.407***	−0.097
	(2.738)	(0.074)	(0.084)
Complementarity	3.323	0.008	−0.012
	(2.726)	(0.082)	(0.081)
Num.Obs.	929	929	930
RMSE	33.27	0.96	1.04

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.2 Partisanship

Democrats

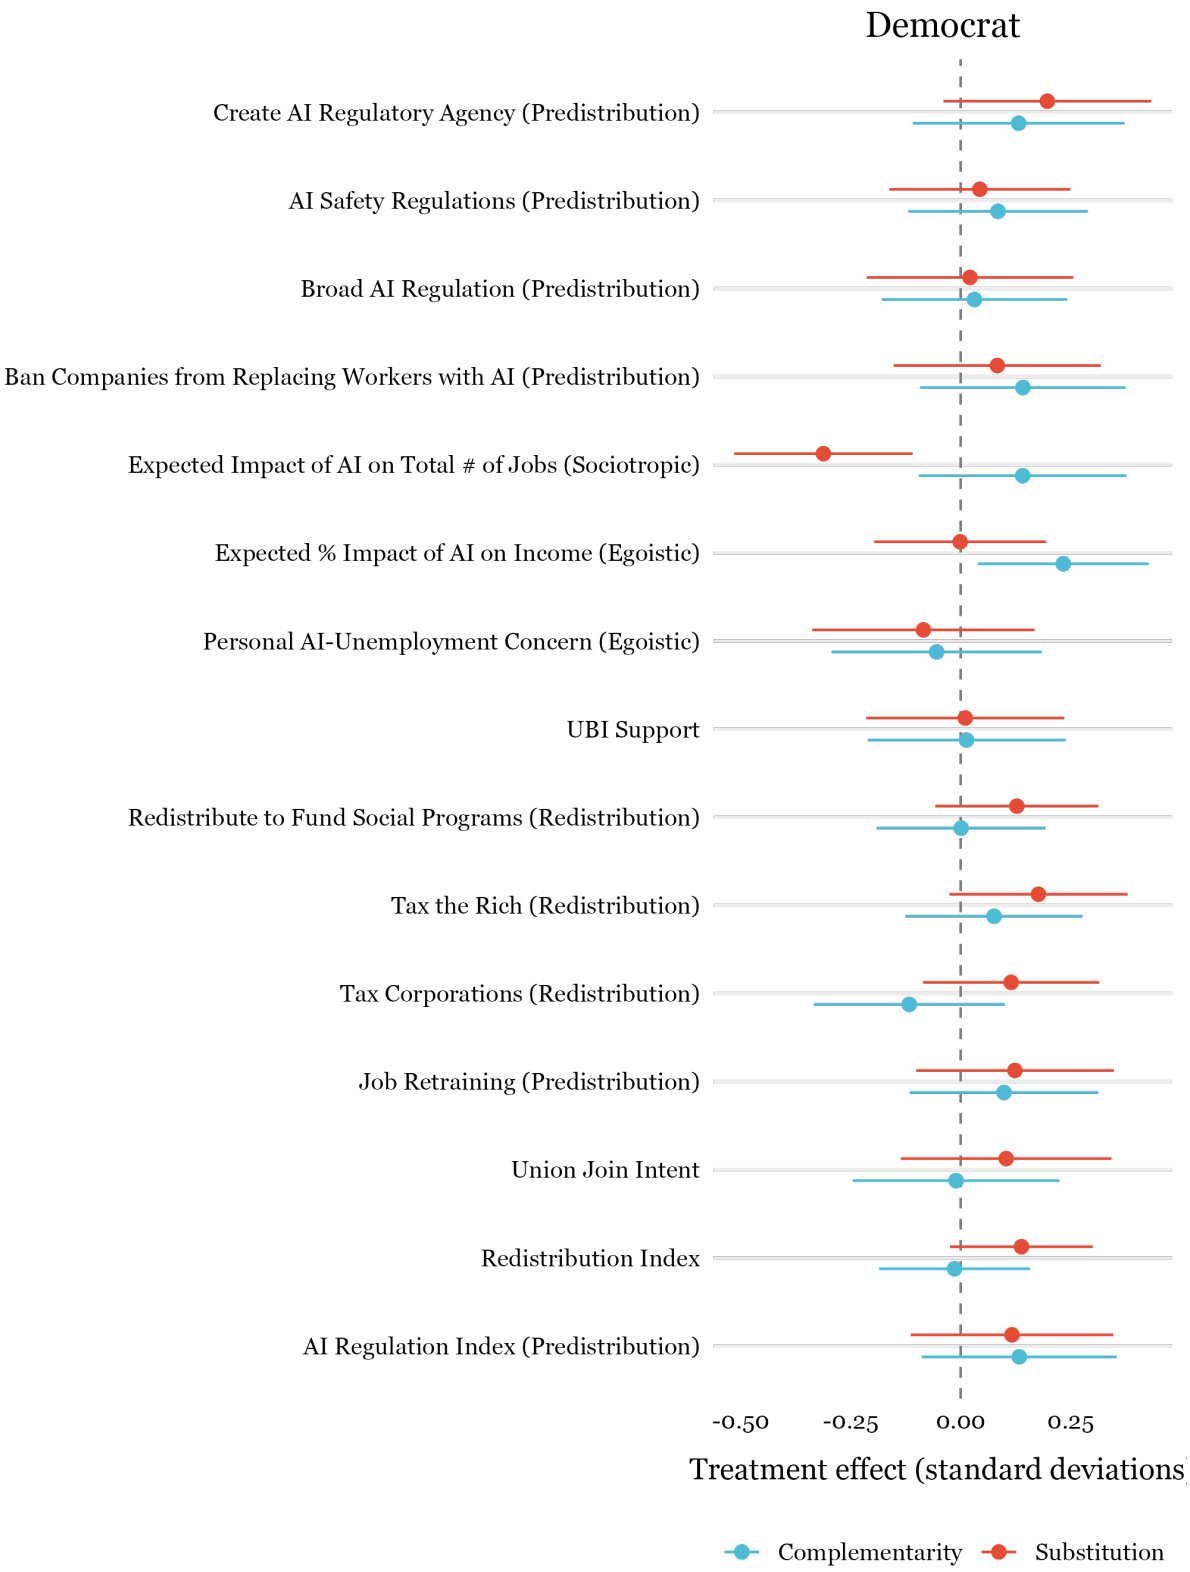


Table C.5: Democrat: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.138+	0.010	0.123	0.117	0.177+	0.115	0.103
	(0.083)	(0.114)	(0.114)	(0.117)	(0.103)	(0.102)	(0.122)
Complementarity	−0.014	0.013	0.099	0.133	0.076	−0.117	−0.010
	(0.087)	(0.114)	(0.109)	(0.113)	(0.102)	(0.110)	(0.119)
Num.Obs.	412	405	411	410	411	411	392
RMSE	0.70	0.93	0.93	0.93	0.85	0.88	0.97

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.6: Democrat: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.021	0.083	0.044	0.197
	(0.119)	(0.120)	(0.105)	(0.120)
Complementarity	0.031	0.141	0.085	0.132
	(0.107)	(0.119)	(0.104)	(0.122)
Num.Obs.	404	410	410	409
RMSE	0.90	0.98	0.86	0.98

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.7: Democrat: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−0.041	−0.312**	−0.085
	(3.655)	(0.103)	(0.129)
Complementarity	8.577*	0.141	−0.055
	(3.627)	(0.120)	(0.122)
Num.Obs.	412	411	412
RMSE	29.26	0.92	1.04

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Republicans

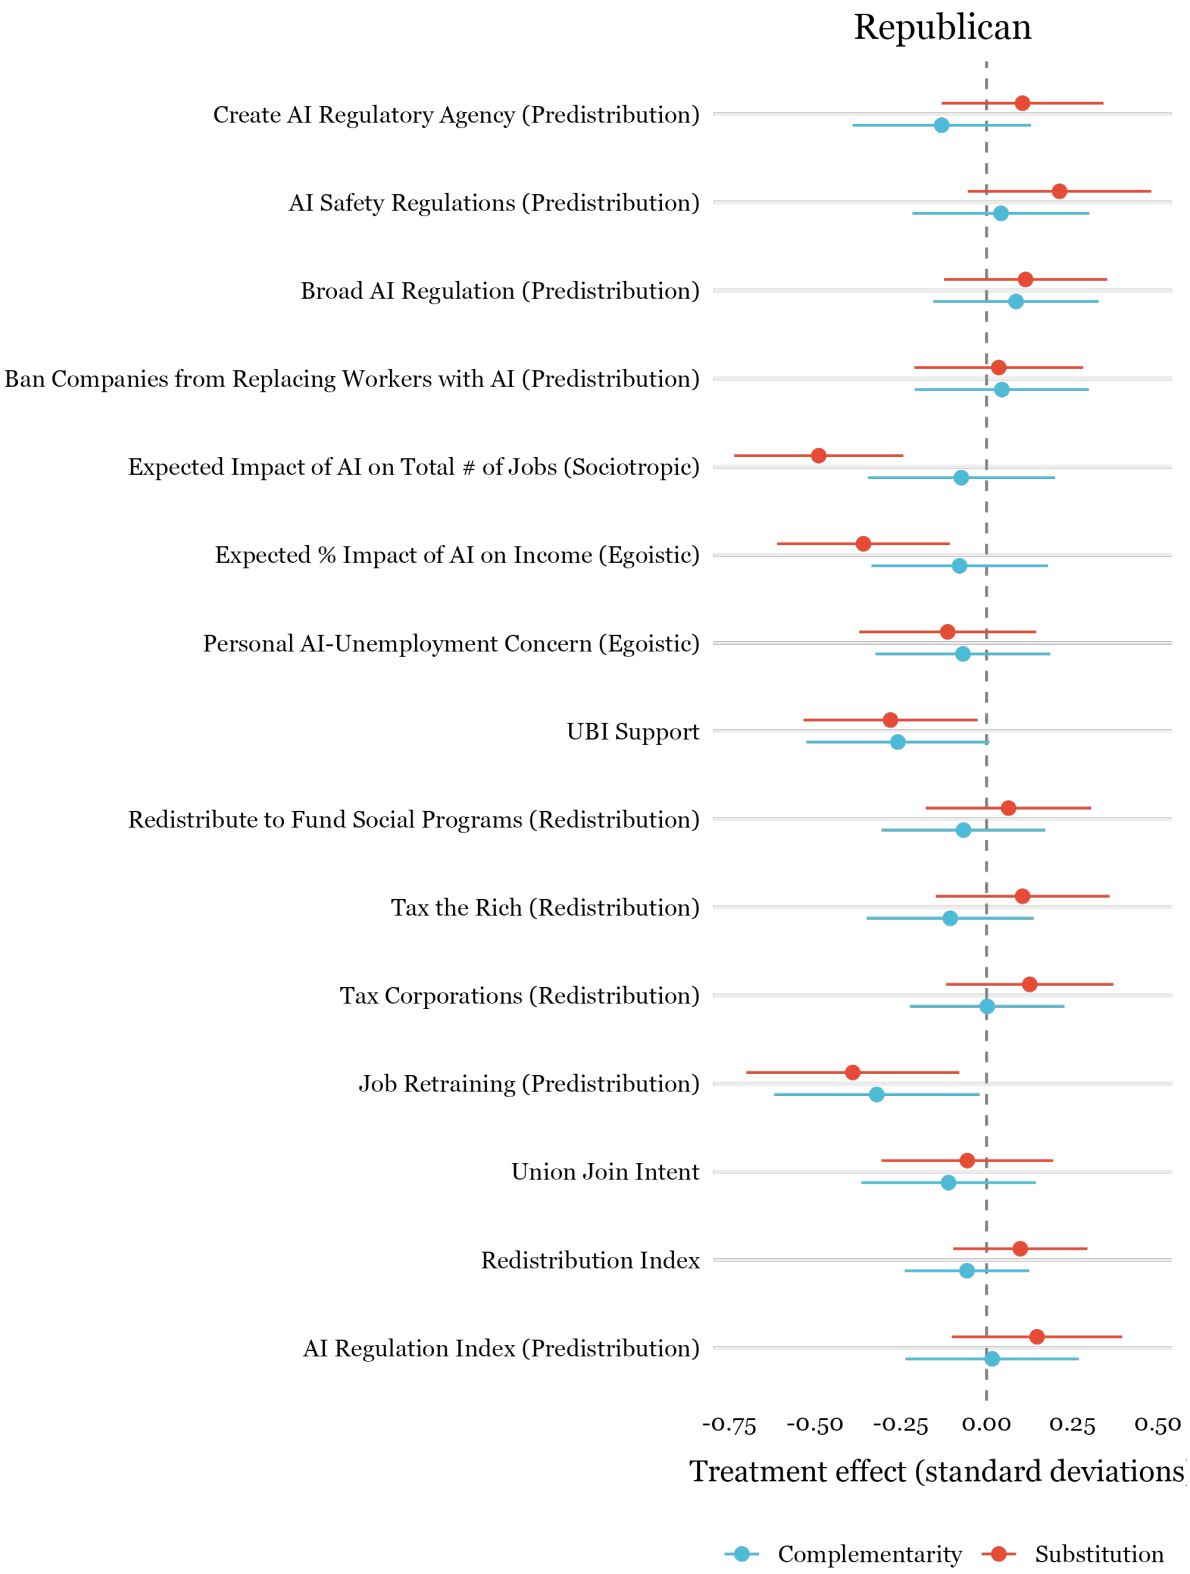


Table C.8: Republican: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.098	−0.280*	−0.390*	0.147	0.105	0.126	−0.056
	(0.099)	(0.129)	(0.158)	(0.126)	(0.129)	(0.124)	(0.127)
Complementarity	−0.057	−0.259+	−0.320*	0.016	−0.106	0.002	−0.111
	(0.092)	(0.136)	(0.152)	(0.128)	(0.124)	(0.114)	(0.129)
Num.Obs.	356	353	355	356	356	356	349
RMSE	0.74	1.00	1.22	0.99	0.97	0.93	0.98

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.9: Republican: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.114	0.035	0.213	0.105
	(0.121)	(0.125)	(0.136)	(0.120)
Complementarity	0.086	0.045	0.042	−0.131
	(0.122)	(0.129)	(0.131)	(0.132)
Num.Obs.	353	356	356	356
RMSE	0.92	0.98	1.03	0.99

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.10: Republican: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−13.195**	−0.489***	−0.113
	(4.705)	(0.125)	(0.131)
Complementarity	−2.900	−0.074	−0.069
	(4.801)	(0.138)	(0.129)
Num.Obs.	356	356	356
RMSE	36.14	1.00	1.01

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Independent and Other

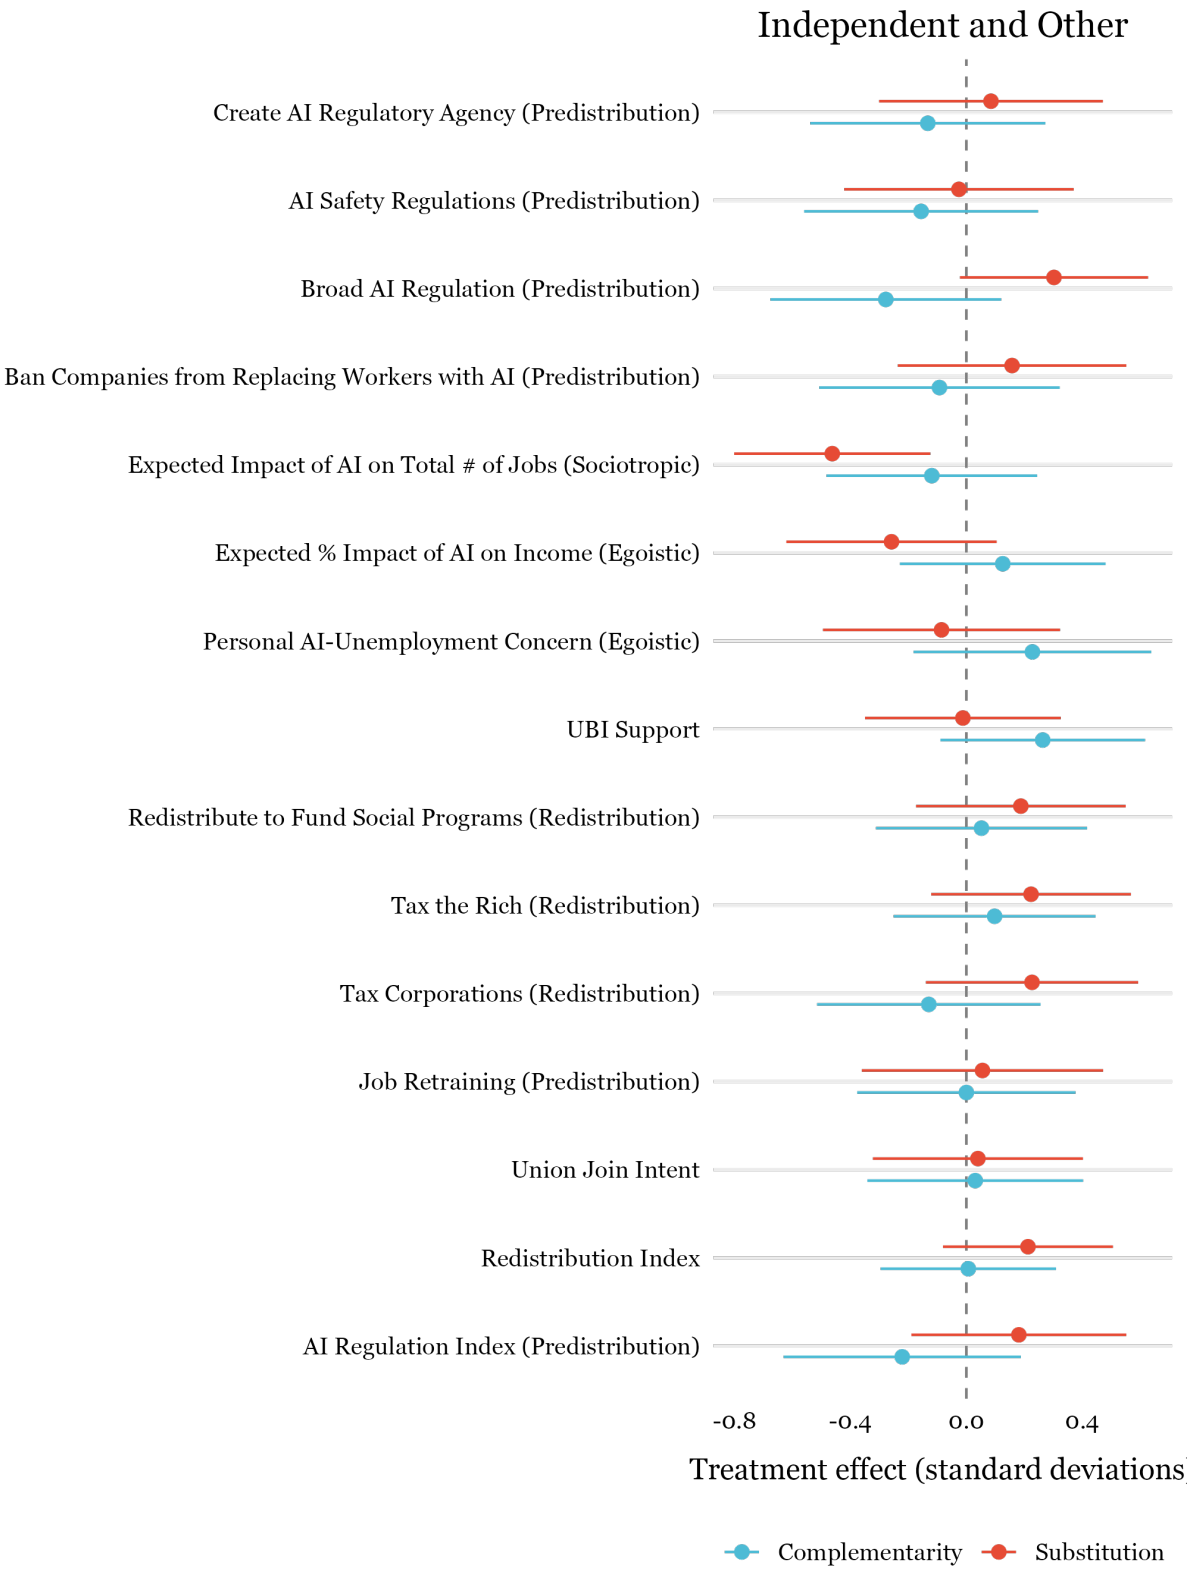


Table C.11: Independent and Other: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.213	−0.012	0.056	0.181	0.223	0.227	0.040
	(0.148)	(0.171)	(0.211)	(0.188)	(0.175)	(0.185)	(0.183)
Complementarity	0.007	0.264	0.000	−0.221	0.097	−0.130	0.031
	(0.153)	(0.179)	(0.191)	(0.207)	(0.176)	(0.195)	(0.188)
Num.Obs.	161	159	161	161	161	161	157
RMSE	0.77	0.89	1.04	1.01	0.88	0.96	0.95

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.12: Independent and Other: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.302+	0.158	−0.026	0.085
	(0.164)	(0.200)	(0.201)	(0.196)
Complementarity	−0.278	−0.093	−0.156	−0.133
	(0.202)	(0.210)	(0.204)	(0.205)
Num.Obs.	160	161	161	161
RMSE	0.95	1.06	1.03	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.13: Independent and Other: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−9.493	−0.463**	−0.086
	(6.756)	(0.171)	(0.207)
Complementarity	4.618	−0.119	0.228
	(6.607)	(0.184)	(0.208)
Num.Obs.	161	162	162
RMSE	33.73	0.89	1.08

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.3 Education

No Degree

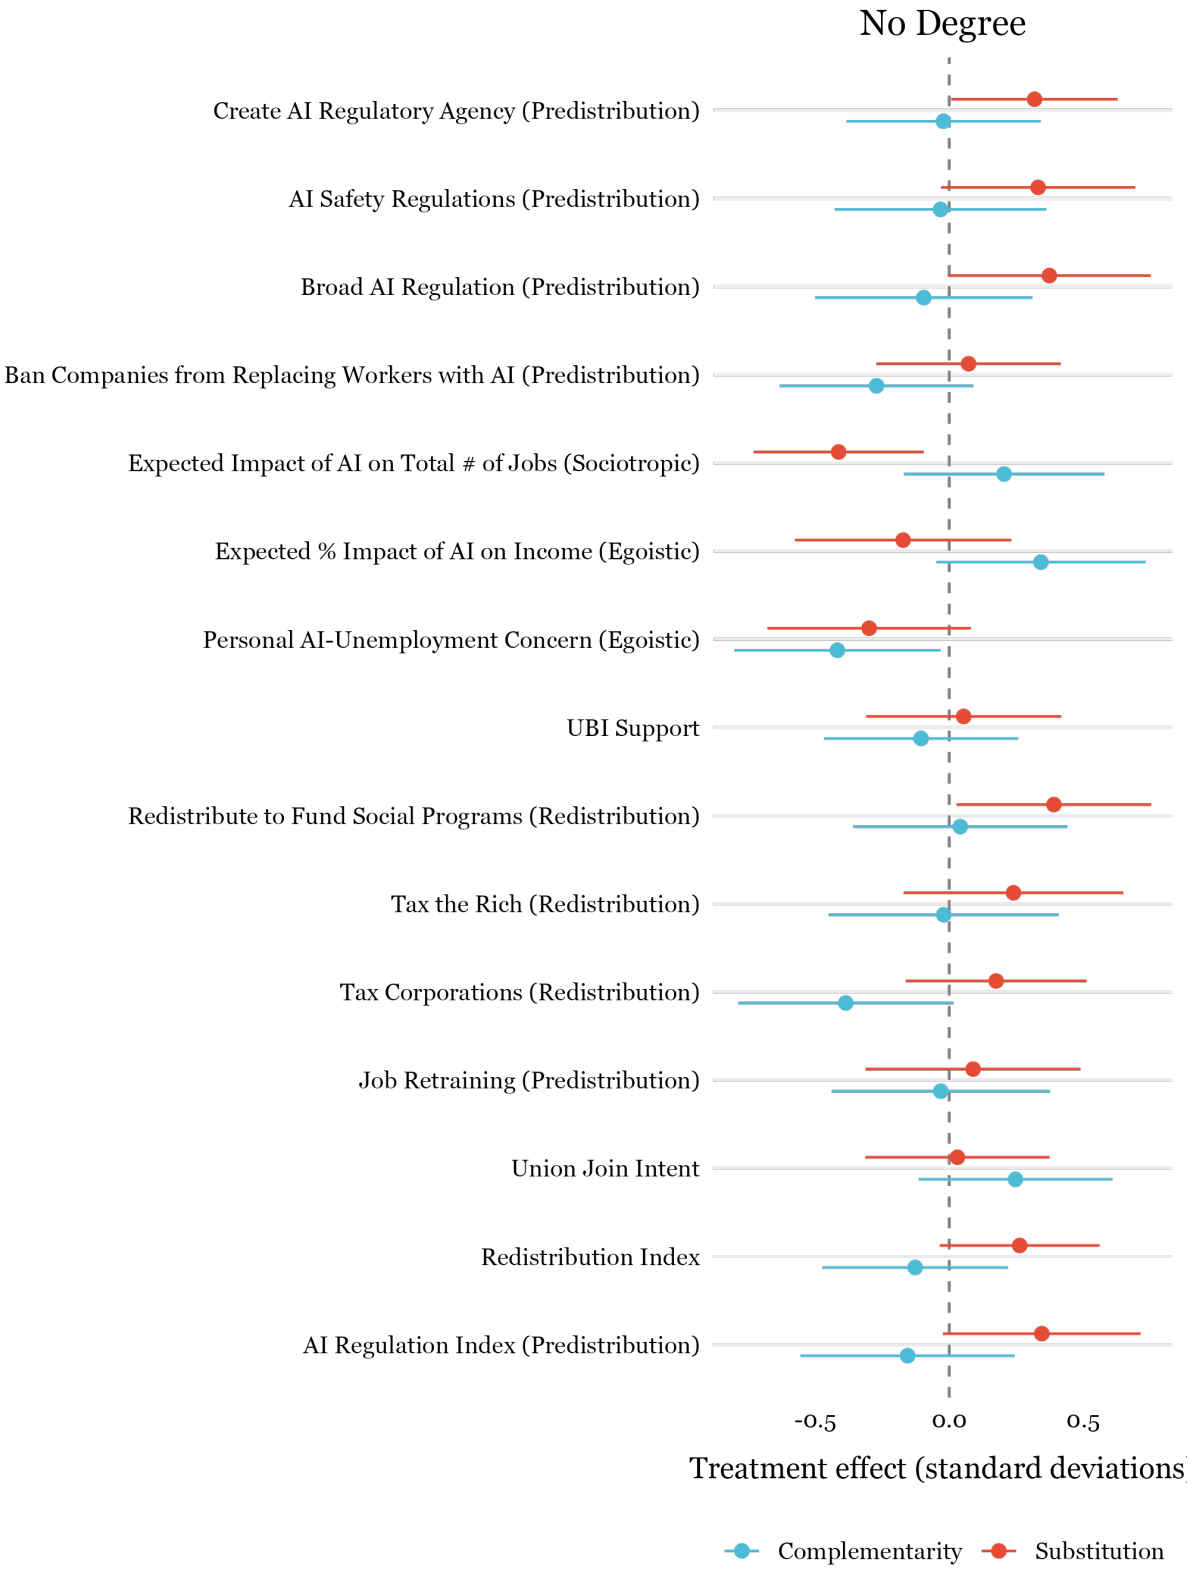


Table C.14: No Degree: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.263+	0.054	0.089	0.346+	0.240	0.175	0.031
	(0.151)	(0.184)	(0.203)	(0.187)	(0.208)	(0.171)	(0.174)
Complementarity	−0.127	−0.105	−0.031	−0.155	−0.020	−0.386+	0.247
	(0.176)	(0.184)	(0.207)	(0.203)	(0.218)	(0.204)	(0.183)
Num.Obs.	178	177	178	177	178	177	175
RMSE	0.88	0.99	1.12	1.03	1.12	1.02	0.97

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.15: No Degree: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.374+	0.072	0.332+	0.319*
	(0.192)	(0.174)	(0.183)	(0.157)
Complementarity	−0.095	−0.271	−0.032	−0.022
	(0.206)	(0.183)	(0.200)	(0.183)
Num.Obs.	177	177	177	177
RMSE	1.05	0.95	1.02	0.91

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.16: No Degree: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−6.319	−0.413*	−0.299
	(7.521)	(0.161)	(0.193)
Complementarity	12.587+	0.205	−0.417*
	(7.280)	(0.189)	(0.195)
Num.Obs.	178	178	178
RMSE	37.96	0.94	1.05

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Degree

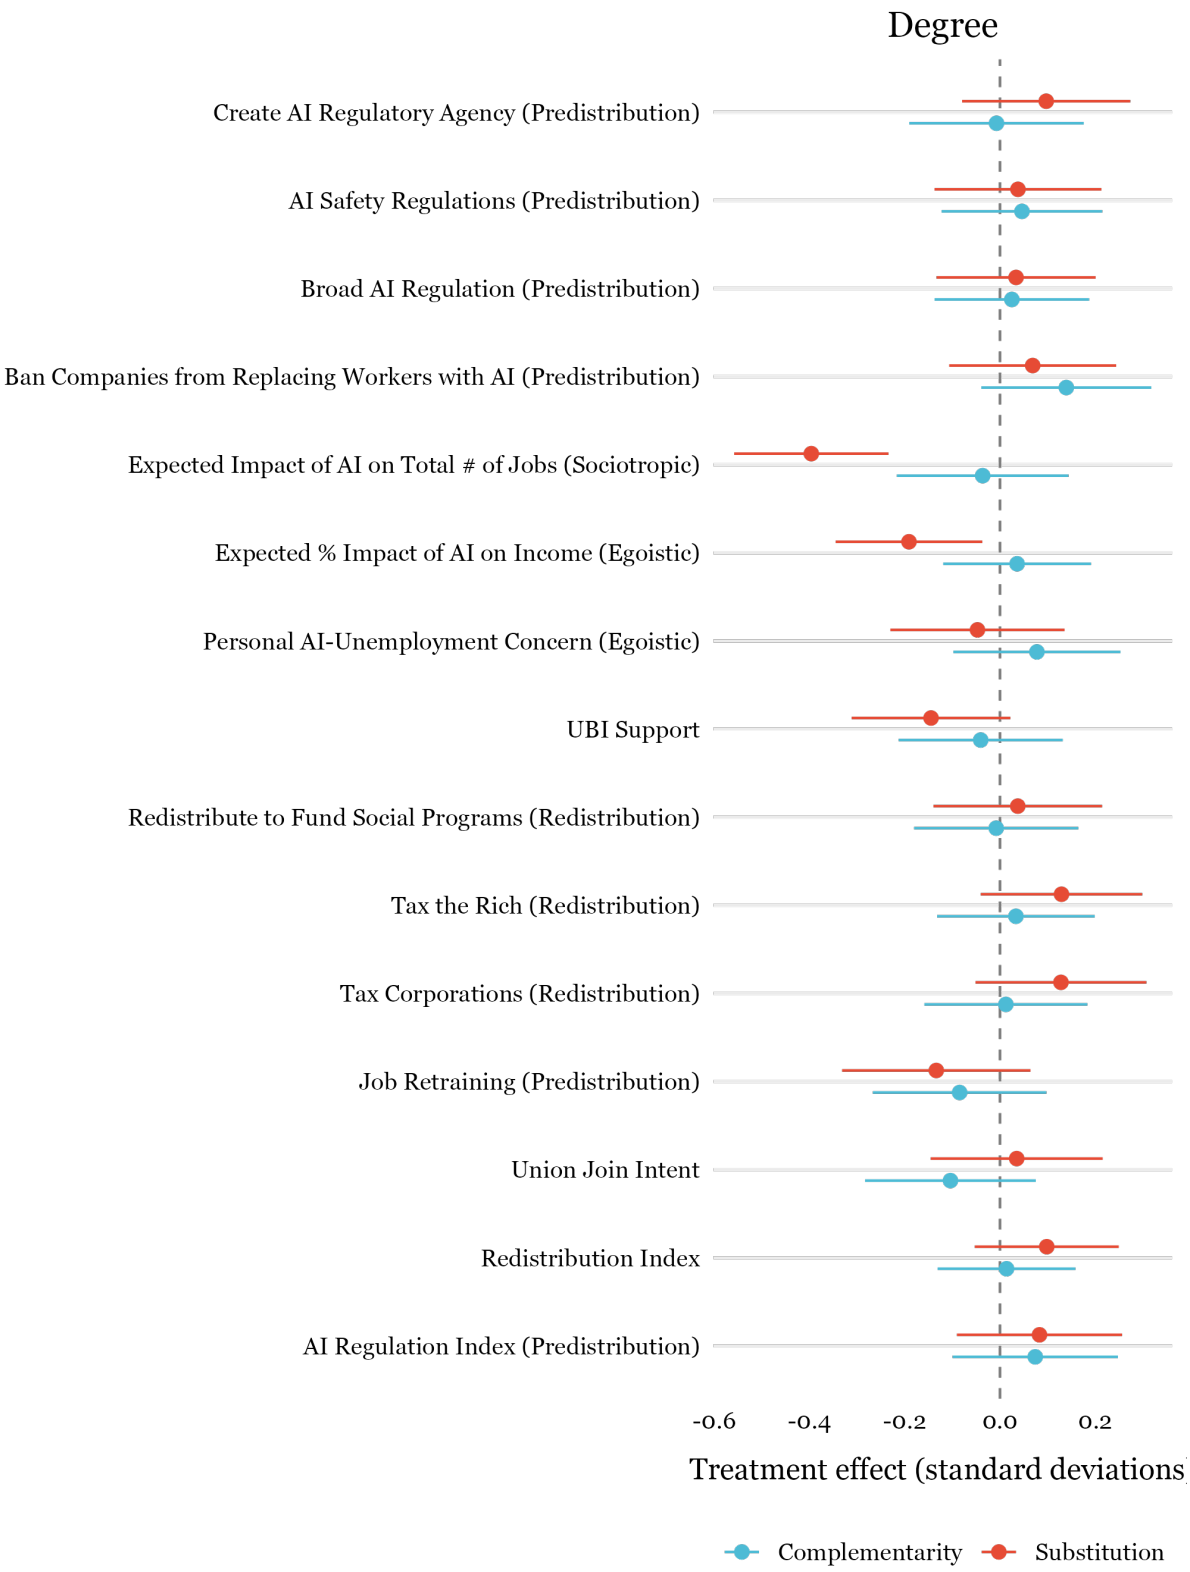

Degree
Create AI Regulatory Agency (Predistribution)
AI Safety Regulations (Predistribution)
Broad AI Regulation (Predistribution)
Ban Companies from Replacing Workers with AI (Predistribution)
Expected Impact of AI on Total # of Jobs (Sociotropic)
Expected % Impact of AI on Income (Egoistic)
Personal AI-Unemployment Concern (Egoistic)
UBI Support
Redistribute to Fund Social Programs (Redistribution)
Tax the Rich (Redistribution)
Tax Corporations (Redistribution)
Job Retraining (Predistribution)
Union Join Intent
Redistribution Index
AI Regulation Index (Predistribution)
-0.6
-0.4
-0.2
0.0
0.2
Treatment effect (standard deviations
Complementarity
Substitution

Table C.17: Degree: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.098	−0.145+	−0.134	0.083	0.129	0.128	0.035
	(0.077)	(0.085)	(0.101)	(0.088)	(0.086)	(0.091)	(0.092)
Complementarity	0.014	−0.041	−0.085	0.074	0.033	0.012	−0.104
	(0.074)	(0.088)	(0.093)	(0.088)	(0.084)	(0.087)	(0.091)
Num.Obs.	747	736	745	746	746	747	719
RMSE	0.84	0.94	1.11	0.99	0.96	0.99	1.01

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.18: Degree: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.034	0.069	0.038	0.097
	(0.085)	(0.089)	(0.089)	(0.090)
Complementarity	0.025	0.139	0.046	−0.007
	(0.083)	(0.091)	(0.086)	(0.093)
Num.Obs.	736	746	746	745
RMSE	0.92	1.01	0.98	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.19: Degree: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−7.025*	−0.396***	−0.047
	(2.883)	(0.083)	(0.093)
Complementarity	1.321	−0.036	0.077
	(2.906)	(0.092)	(0.089)
Num.Obs.	747	747	748
RMSE	31.99	0.96	1.03

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.4 Age

Under 40

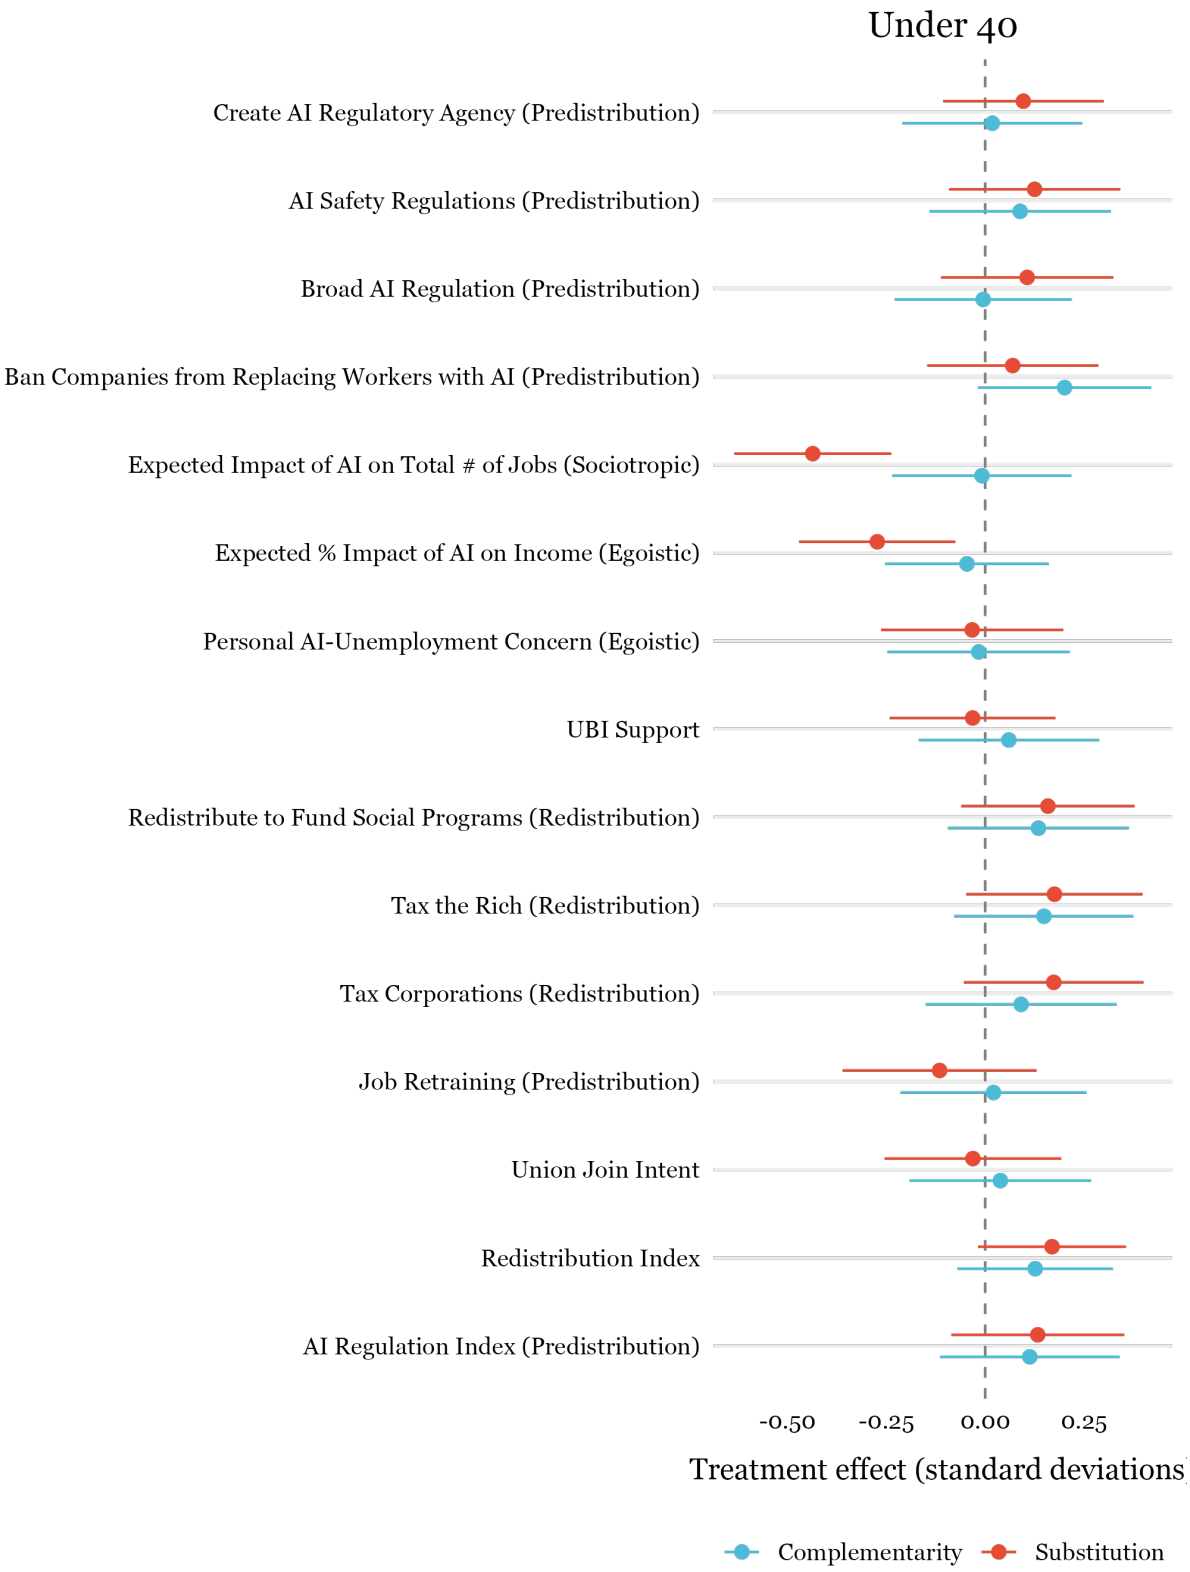


Table C.20: Under 40: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.169+	−0.032	−0.115	0.133	0.175	0.173	−0.031
	(0.095)	(0.107)	(0.125)	(0.111)	(0.114)	(0.116)	(0.114)
Complementarity	0.126	0.060	0.021	0.113	0.149	0.091	0.038
	(0.100)	(0.116)	(0.120)	(0.116)	(0.115)	(0.123)	(0.117)
Num.Obs.	467	459	467	467	466	467	449
RMSE	0.87	0.97	1.11	1.01	1.02	1.06	1.01

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.21: Under 40: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.106	0.070	0.125	0.097
	(0.111)	(0.110)	(0.110)	(0.103)
Complementarity	−0.005	0.201+	0.088	0.018
	(0.114)	(0.112)	(0.117)	(0.116)
Num.Obs.	462	467	467	467
RMSE	0.98	0.98	1.01	0.98

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.22: Under 40: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−10.034**	−0.436***	−0.033
	(3.702)	(0.101)	(0.117)
Complementarity	−1.692	−0.008	−0.016
	(3.874)	(0.115)	(0.117)
Num.Obs.	468	468	468
RMSE	32.88	0.95	1.05

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

40 and Over

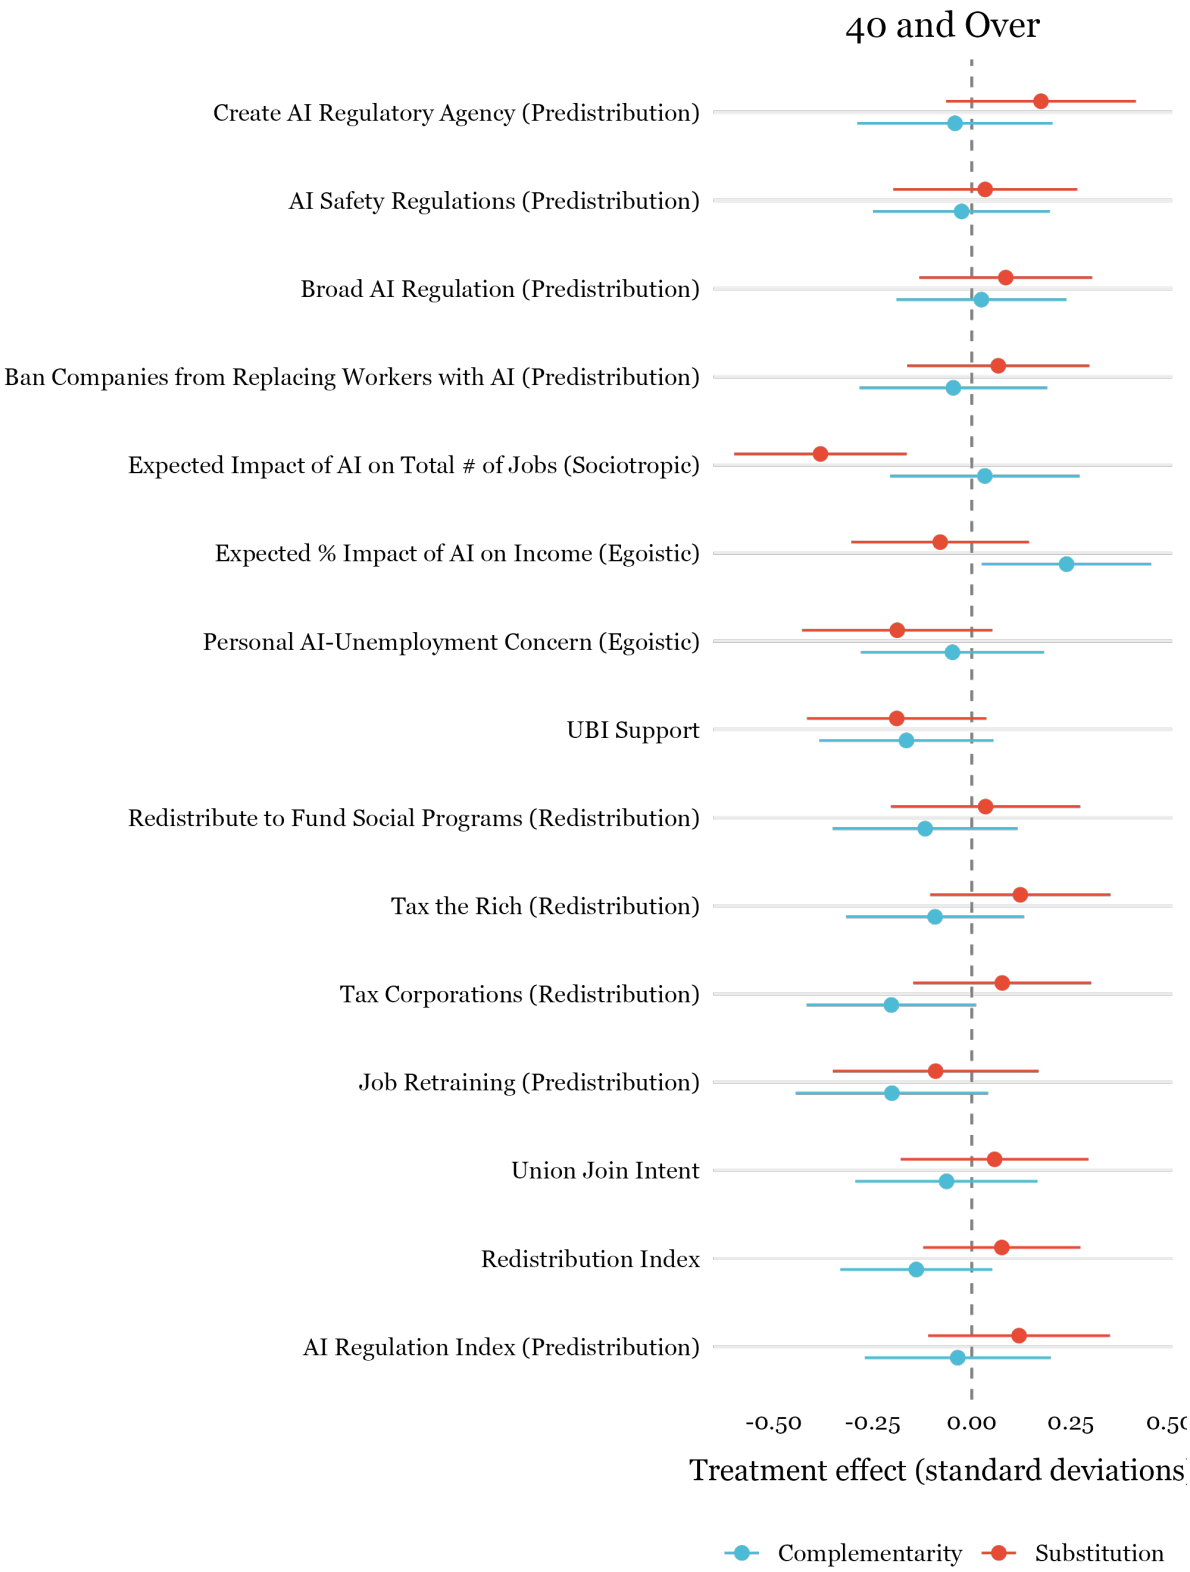

40 and Over
Create AI Regulatory Agency (Predistribution)
AI Safety Regulations (Predistribution)
Broad AI Regulation (Predistribution)
Ban Companies from Replacing Workers with AI (Predistribution)
Expected Impact of AI on Total # of Jobs (Sociotropic)
Expected % Impact of AI on Income (Egoistic)
Personal AI-Unemployment Concern (Egoistic)
UBI Support
Redistribute to Fund Social Programs (Redistribution)
Tax the Rich (Redistribution)
Tax Corporations (Redistribution)
Job Retraining (Predistribution)
Union Join Intent
Redistribution Index
AI Regulation Index (Predistribution)
-0.50
-0.25
0.00
0.25
Treatment effect (standard deviations
Complementarity
Substitution

Table C.23: 40 and Over: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.076	−0.189	−0.091	0.119	0.123	0.077	0.058
	(0.101)	(0.115)	(0.132)	(0.117)	(0.116)	(0.114)	(0.121)
Complementarity	−0.140	−0.165	−0.202	−0.036	−0.093	−0.203+	−0.064
	(0.098)	(0.112)	(0.124)	(0.119)	(0.114)	(0.109)	(0.117)
Num.Obs.	440	436	438	439	440	439	427
RMSE	0.84	0.94	1.11	1.00	0.97	0.95	0.99

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.24: 40 and Over: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.086	0.067	0.034	0.175
	(0.111)	(0.117)	(0.118)	(0.122)
Complementarity	0.024	−0.047	−0.026	−0.042
	(0.109)	(0.121)	(0.113)	(0.125)
Num.Obs.	434	439	439	438
RMSE	0.92	1.02	0.98	1.03

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.25: 40 and Over: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−2.932	−0.382***	−0.188
	(4.193)	(0.111)	(0.122)
Complementarity	8.797*	0.033	−0.049
	(4.001)	(0.122)	(0.118)
Num.Obs.	439	439	440
RMSE	34.00	0.97	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.5 Sex

Male

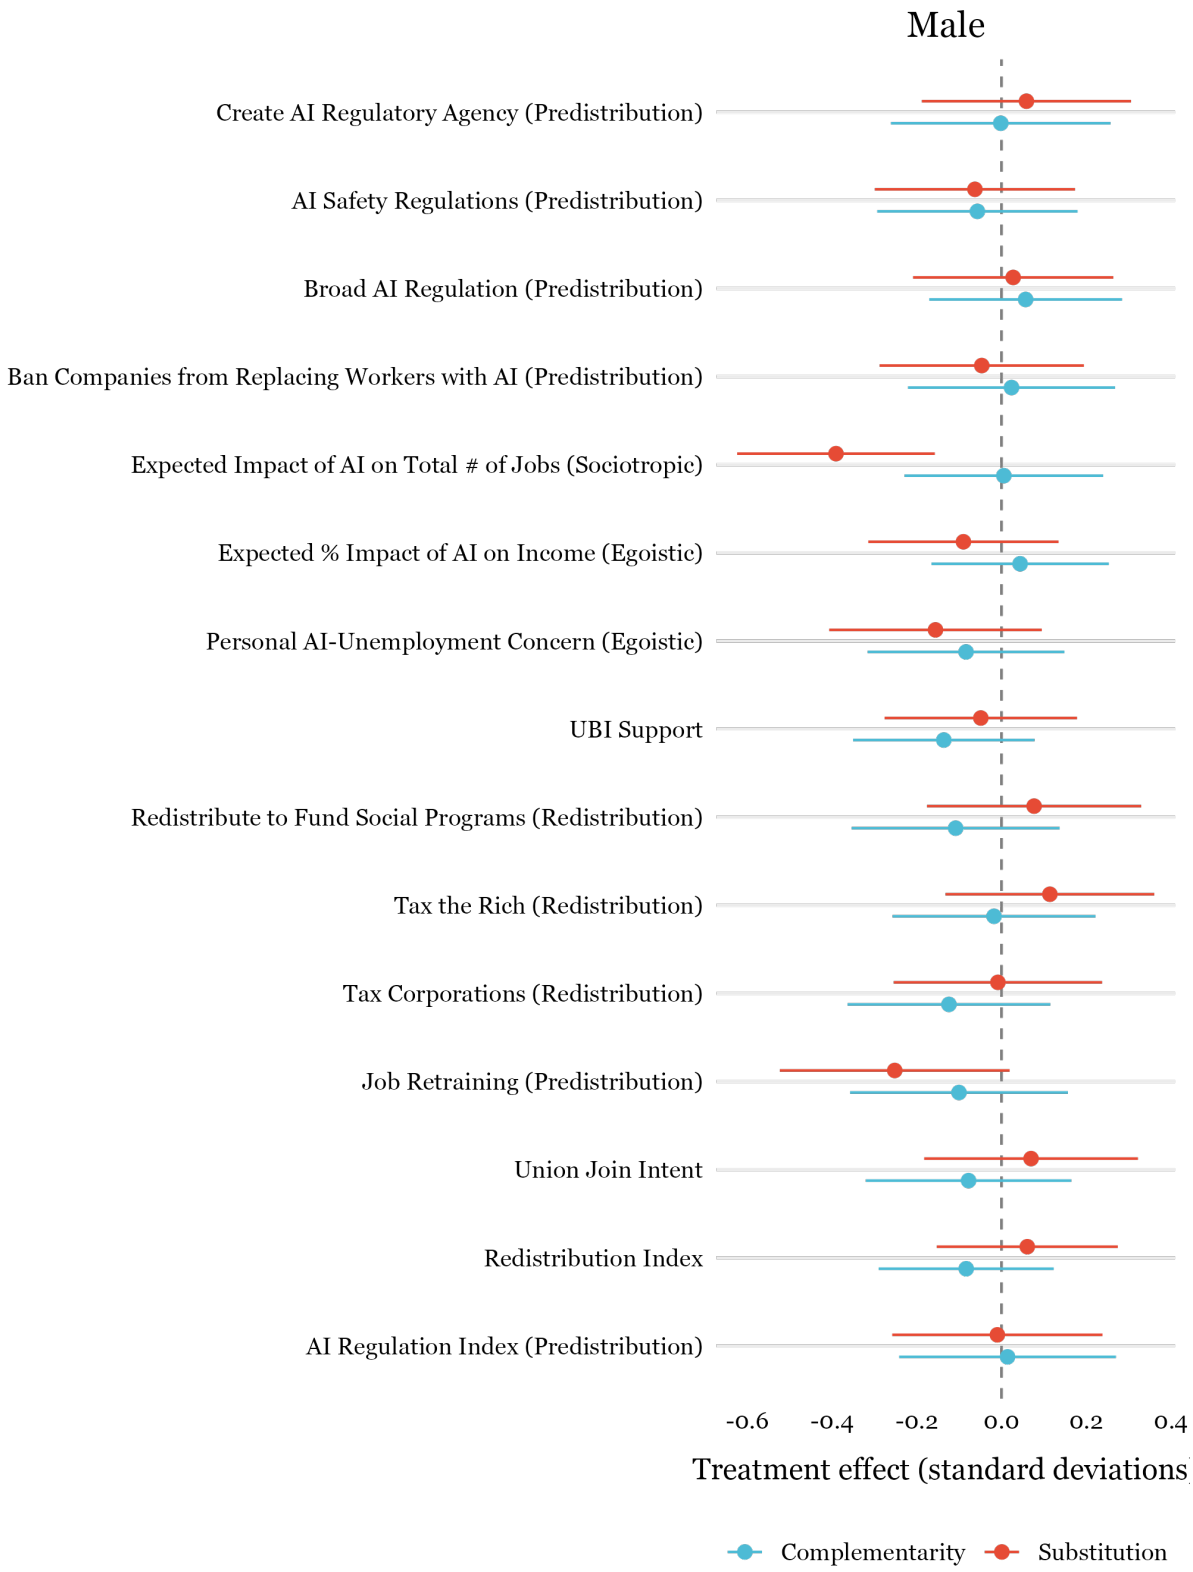

Male
Create AI Regulatory Agency (Predistribution)
AI Safety Regulations (Predistribution)
Broad AI Regulation (Predistribution)
Ban Companies from Replacing Workers with AI (Predistribution)
Expected Impact of AI on Total # of Jobs (Sociotropic)
Expected % Impact of AI on Income (Egoistic)
Personal AI-Unemployment Concern (Egoistic)
UBI Support
Redistribute to Fund Social Programs (Redistribution)
Tax the Rich (Redistribution)
Tax Corporations (Redistribution)
Job Retraining (Predistribution)
Union Join Intent
Redistribution Index
AI Regulation Index (Predistribution)
-0.6
-0.4
-0.2
0.0
0.2
0.4
Treatment effect (standard deviations
Complementarity
Substitution

Table C.26: Male: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.061	−0.049	−0.252+	−0.010	0.114	−0.009	0.070
	(0.109)	(0.116)	(0.138)	(0.126)	(0.125)	(0.125)	(0.128)
Complementarity	−0.083	−0.136	−0.100	0.014	−0.018	−0.124	−0.078
	(0.105)	(0.109)	(0.131)	(0.130)	(0.122)	(0.122)	(0.123)
Num.Obs.	418	412	417	417	418	418	404
RMSE	0.90	0.90	1.15	1.06	1.03	1.05	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.27: Male: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.028	−0.047	−0.063	0.059
	(0.120)	(0.123)	(0.120)	(0.126)
Complementarity	0.057	0.024	−0.057	−0.002
	(0.116)	(0.124)	(0.120)	(0.132)
Num.Obs.	414	417	417	417
RMSE	0.96	1.02	1.01	1.06

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.28: Male: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−3.304	−0.391**	−0.156
	(4.200)	(0.119)	(0.128)
Complementarity	1.617	0.005	−0.084
	(3.915)	(0.119)	(0.118)
Num.Obs.	418	419	419
RMSE	31.98	0.95	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Female

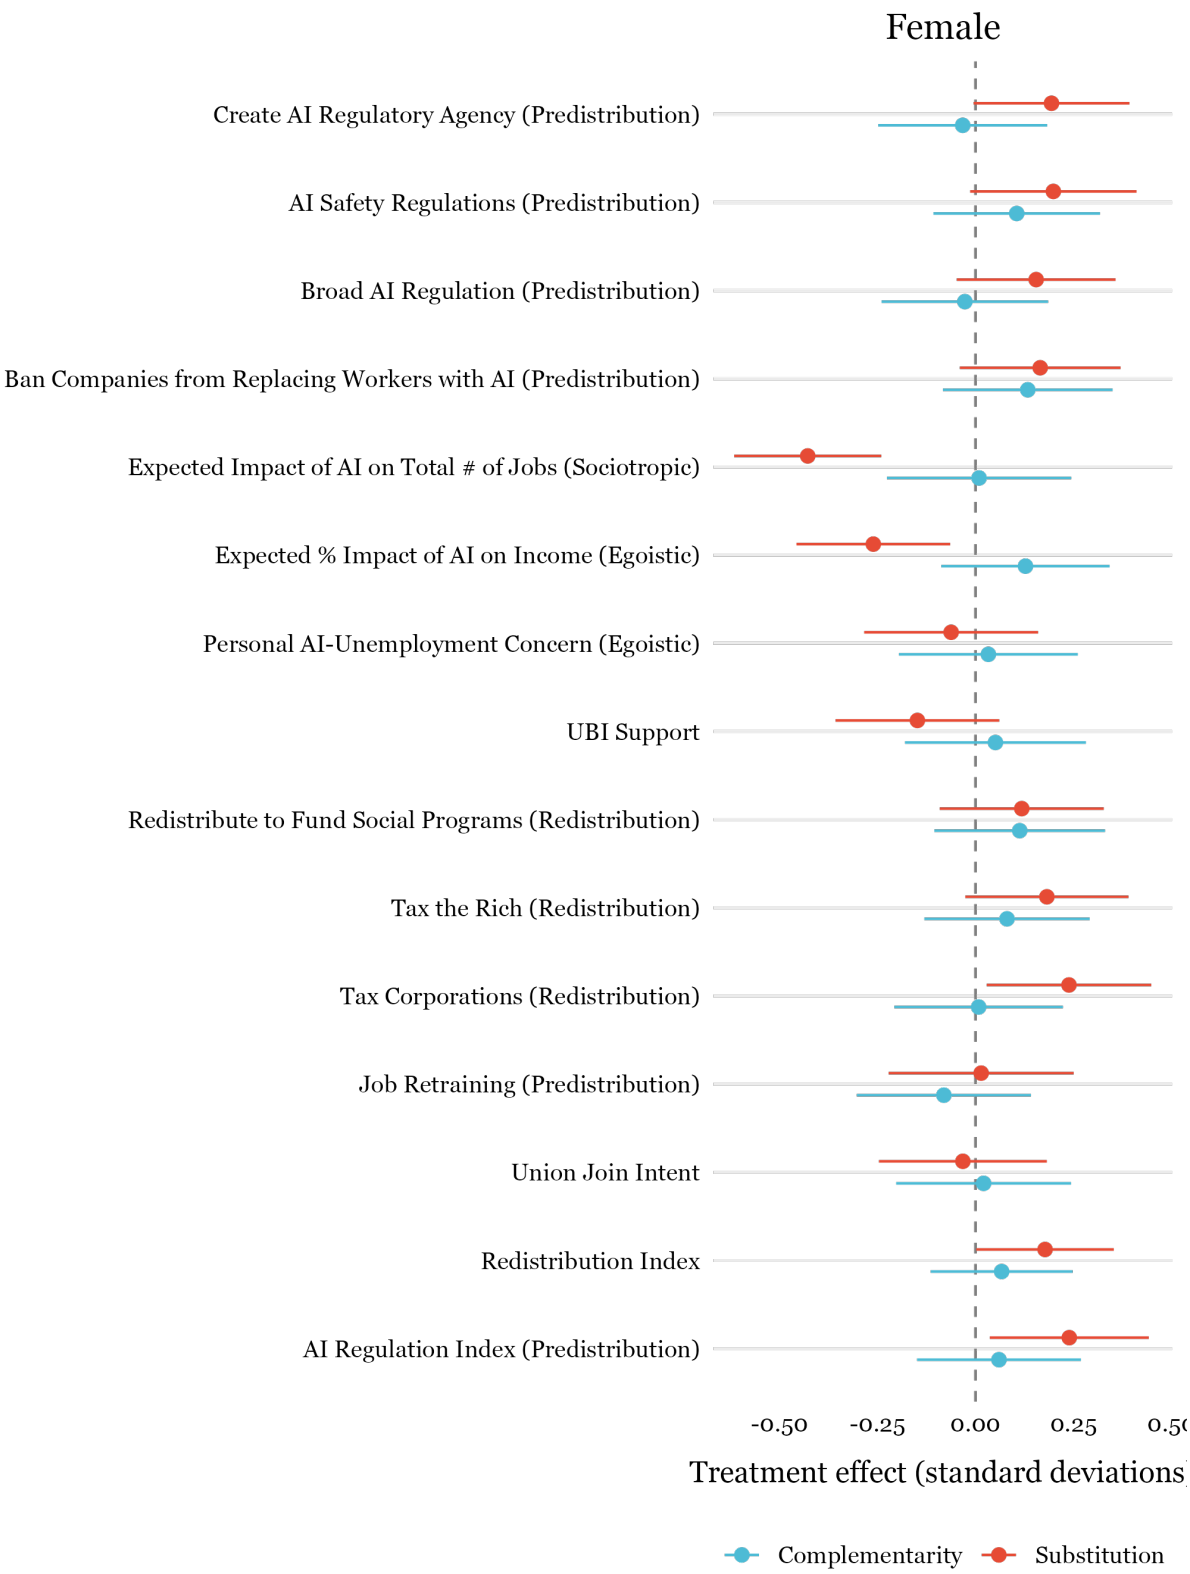

Female
Create AI Regulatory Agency (Predistribution)
AI Safety Regulations (Predistribution)
Broad AI Regulation (Predistribution)
Ban Companies from Replacing Workers with AI (Predistribution)
Expected Impact of AI on Total # of Jobs (Sociotropic)
Expected % Impact of AI on Income (Egoistic)
Personal AI-Unemployment Concern (Egoistic)
UBI Support
Redistribute to Fund Social Programs (Redistribution)
Tax the Rich (Redistribution)
Tax Corporations (Redistribution)
Job Retraining (Predistribution)
Union Join Intent
Redistribution Index
AI Regulation Index (Predistribution)
-0.50
-0.25
0.00
0.25
Treatment effect (standard deviations
Complementarity
Substitution

Table C.29: Female: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.177*	−0.148	0.014	0.239*	0.182+	0.239*	−0.032
	(0.089)	(0.106)	(0.120)	(0.103)	(0.106)	(0.107)	(0.109)
Complementarity	0.067	0.051	−0.081	0.060	0.080	0.008	0.021
	(0.092)	(0.117)	(0.113)	(0.106)	(0.107)	(0.109)	(0.113)
Num.Obs.	489	483	488	489	488	488	472
RMSE	0.81	1.00	1.08	0.95	0.96	0.97	0.99

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.30: Female: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.155	0.165	0.199+	0.194+
	(0.103)	(0.104)	(0.108)	(0.101)
Complementarity	−0.027	0.133	0.105	−0.033
	(0.108)	(0.110)	(0.108)	(0.110)
Num.Obs.	482	489	489	488
RMSE	0.94	0.97	0.97	0.95

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.31: Female: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−9.579**	−0.428***	−0.062
	(3.669)	(0.096)	(0.113)
Complementarity	4.687	0.009	0.033
	(4.017)	(0.120)	(0.116)
Num.Obs.	489	488	489
RMSE	34.41	0.97	1.05

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.6 AI Knowledge

Low AI Knowledge

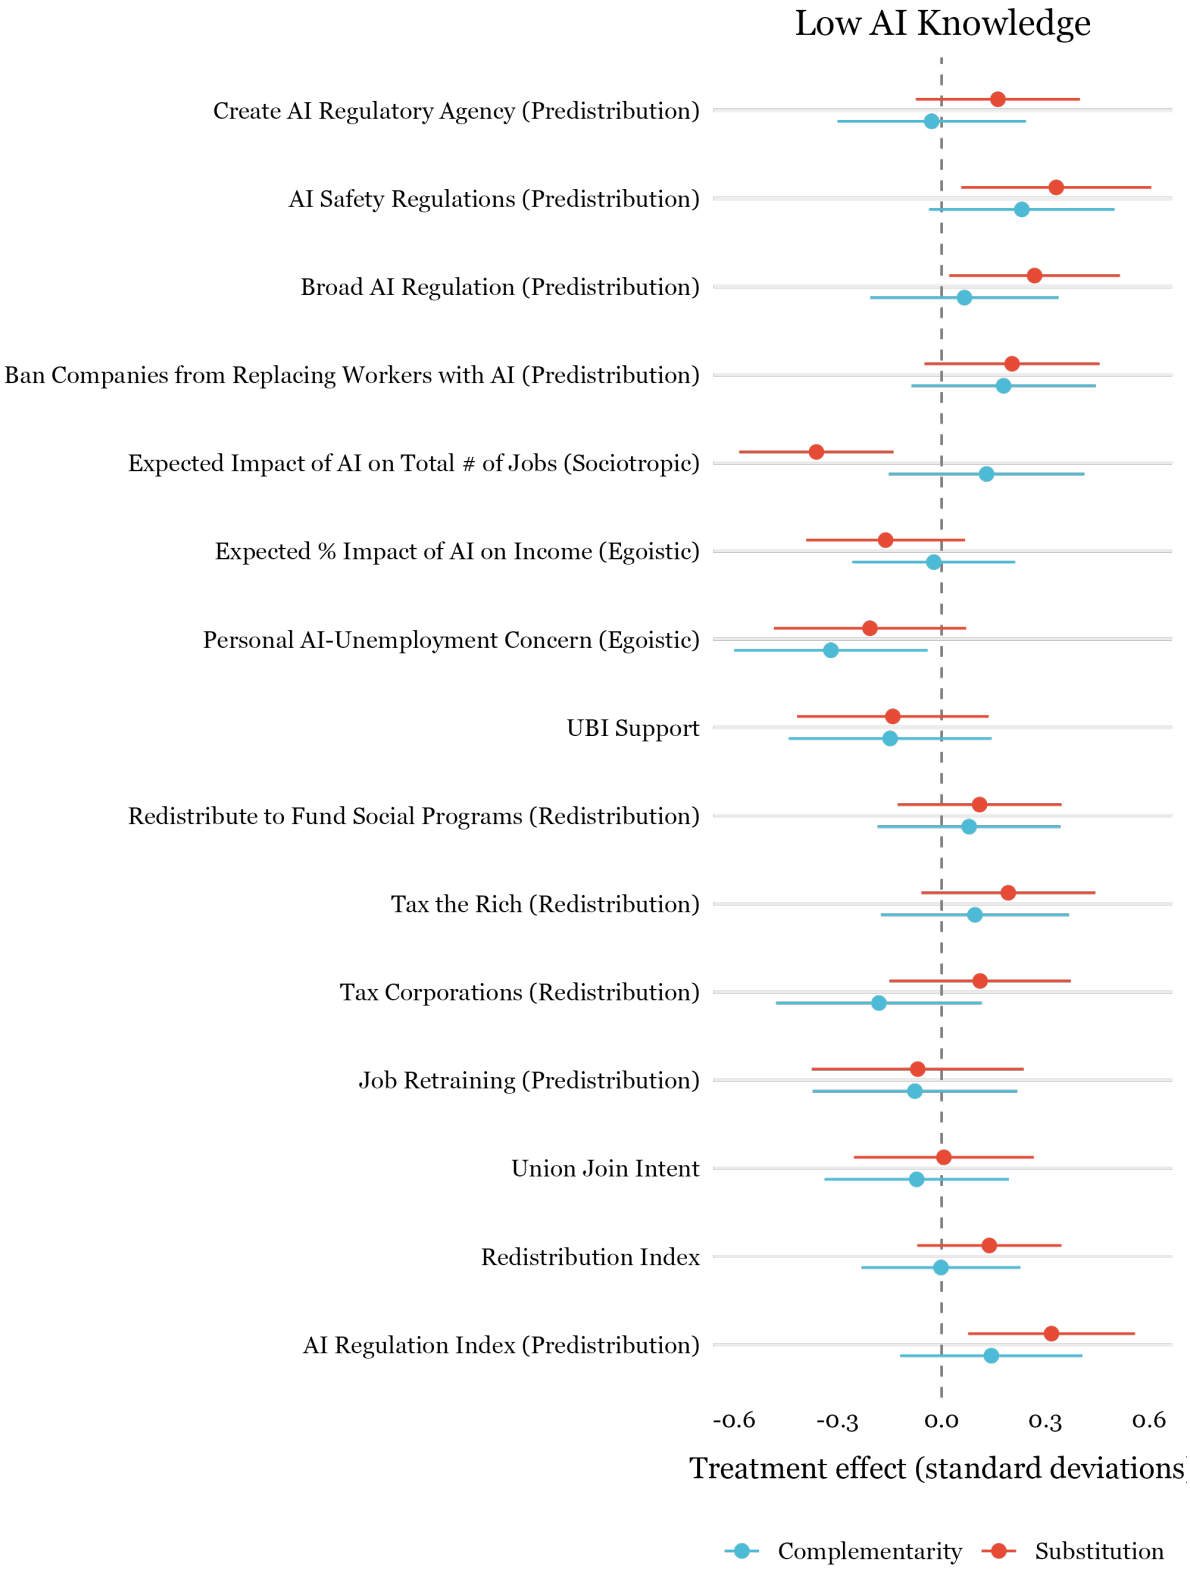


Table C.32: Low AI Knowledge: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.138	−0.141	−0.069	0.318**	0.193	0.111	0.007
	(0.106)	(0.141)	(0.156)	(0.123)	(0.128)	(0.133)	(0.132)
Complementarity	−0.002	−0.149	−0.077	0.144	0.097	−0.181	−0.072
	(0.117)	(0.149)	(0.150)	(0.134)	(0.138)	(0.151)	(0.135)
Num.Obs.	307	304	307	306	307	307	294
RMSE	0.79	1.02	1.12	0.91	0.94	1.02	0.94

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.33: Low AI Knowledge: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.269*	0.204	0.331*	0.163
	(0.125)	(0.129)	(0.140)	(0.121)
Complementarity	0.066	0.179	0.232+	−0.029
	(0.138)	(0.135)	(0.136)	(0.138)
Num.Obs.	301	306	306	306
RMSE	0.93	0.95	0.99	0.92

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.34: Low AI Knowledge: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−5.944	−0.362**	−0.207
	(4.290)	(0.113)	(0.141)
Complementarity	−0.821	0.130	−0.320*
	(4.392)	(0.144)	(0.142)
Num.Obs.	306	307	307
RMSE	30.48	0.90	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

High AI Knowledge

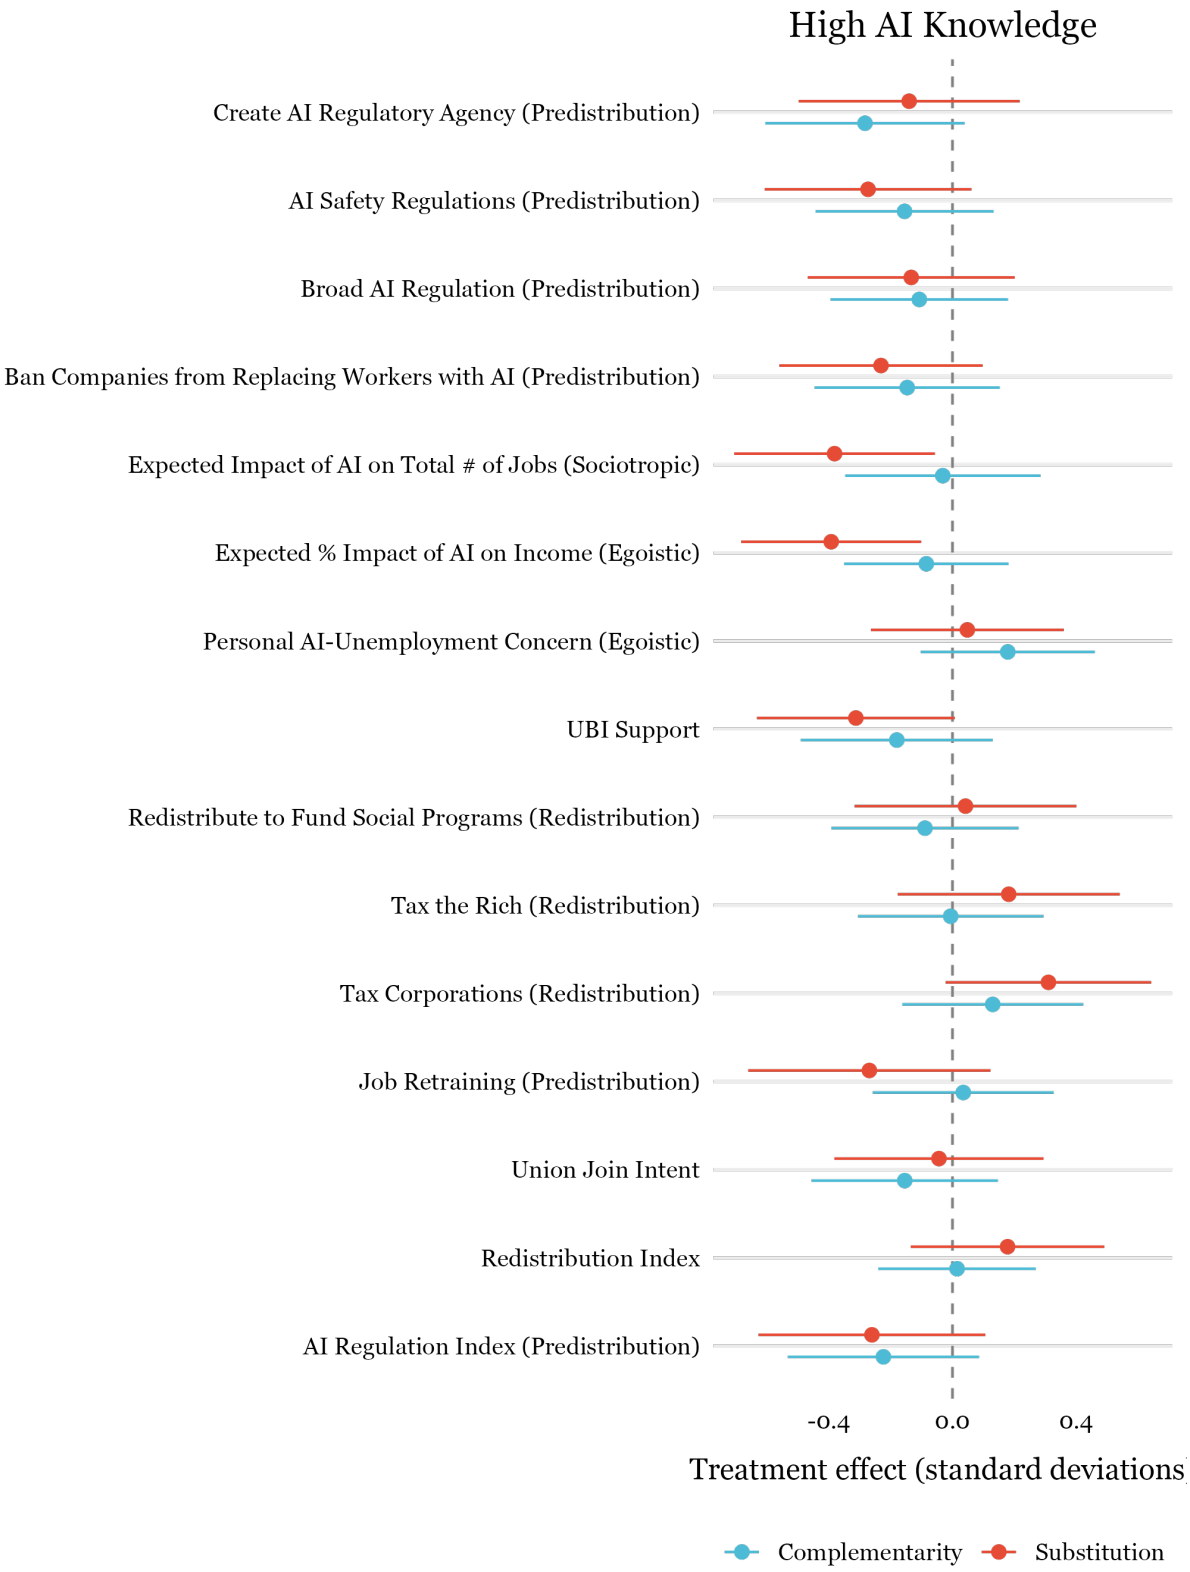


Table C.35: High AI Knowledge: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.178	−0.312+	−0.269	−0.260	0.182	0.310+	−0.044
	(0.159)	(0.162)	(0.199)	(0.186)	(0.182)	(0.169)	(0.172)
Complementarity	0.015	−0.180	0.035	−0.223	−0.006	0.130	−0.154
	(0.129)	(0.157)	(0.148)	(0.157)	(0.152)	(0.148)	(0.153)
Num.Obs.	237	235	236	237	236	237	230
RMSE	0.91	0.99	1.09	1.10	1.06	1.00	1.01

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.36: High AI Knowledge: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	−0.133	−0.231	−0.273	−0.140
	(0.170)	(0.167)	(0.170)	(0.181)
Complementarity	−0.107	−0.146	−0.155	−0.283+
	(0.146)	(0.152)	(0.146)	(0.164)
Num.Obs.	234	237	237	237
RMSE	0.98	1.01	1.01	1.11

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.37: High AI Knowledge: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−14.411**	−0.381*	0.048
	(5.428)	(0.165)	(0.158)
Complementarity	−3.105	−0.031	0.179
	(4.964)	(0.160)	(0.143)
Num.Obs.	238	238	238
RMSE	31.52	1.03	0.96

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.7 AI Self Efficacy

Low AI Self-Efficacy

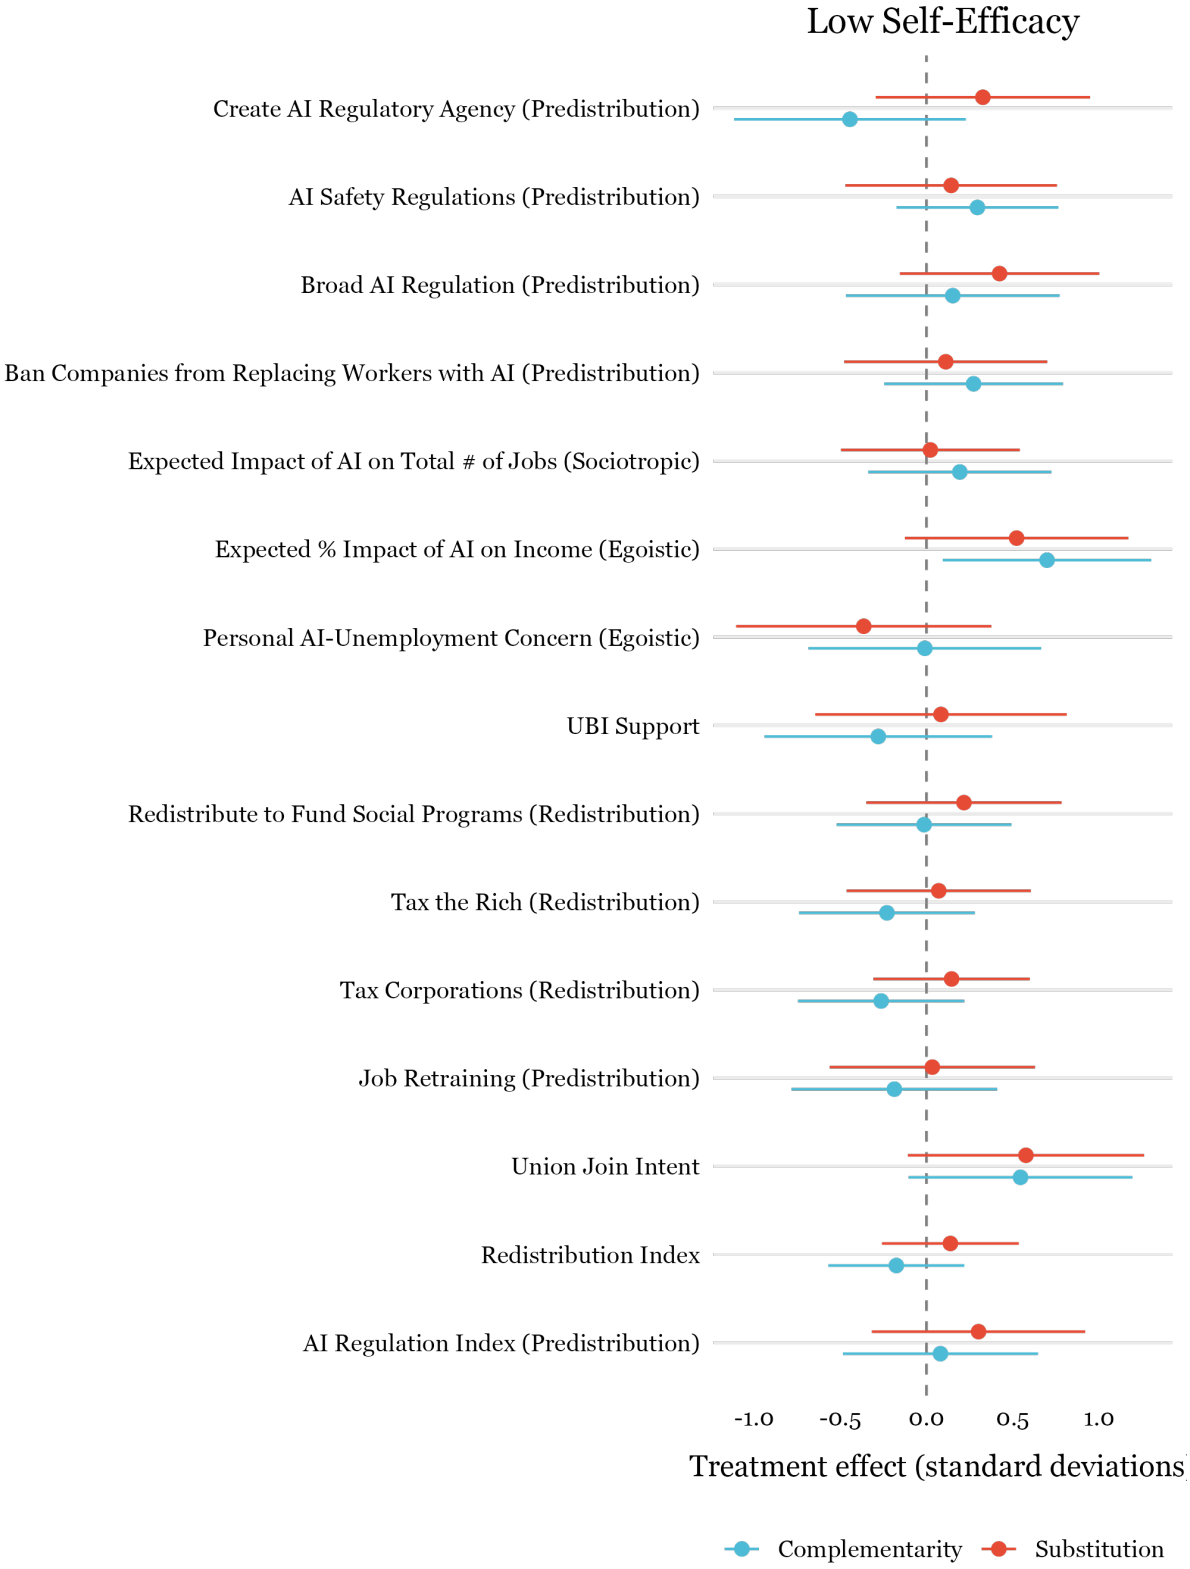


Table C.38: Low AI Self-Efficacy: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.139	0.084	0.034	0.302	0.071	0.146	0.577+
	(0.199)	(0.365)	(0.298)	(0.310)	(0.267)	(0.227)	(0.342)
Complementarity	−0.175	−0.280	−0.186	0.081	−0.229	−0.263	0.545+
	(0.197)	(0.330)	(0.299)	(0.283)	(0.255)	(0.241)	(0.324)
Num.Obs.	67	66	67	67	67	66	61
RMSE	0.69	1.03	1.03	0.93	0.90	0.79	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.39: Low AI Self-Efficacy: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.424	0.112	0.143	0.327
	(0.289)	(0.295)	(0.307)	(0.311)
Complementarity	0.152	0.273	0.295	−0.444
	(0.310)	(0.260)	(0.235)	(0.336)
Num.Obs.	64	67	67	67
RMSE	0.92	0.84	0.89	1.06

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.40: Low AI Self-Efficacy: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	19.208	0.023	−0.364
	(11.925)	(0.260)	(0.370)
Complementarity	25.696*	0.193	−0.009
	(11.123)	(0.266)	(0.338)
Num.Obs.	67	67	67
RMSE	34.82	0.81	1.17

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

High AI Self-Efficacy

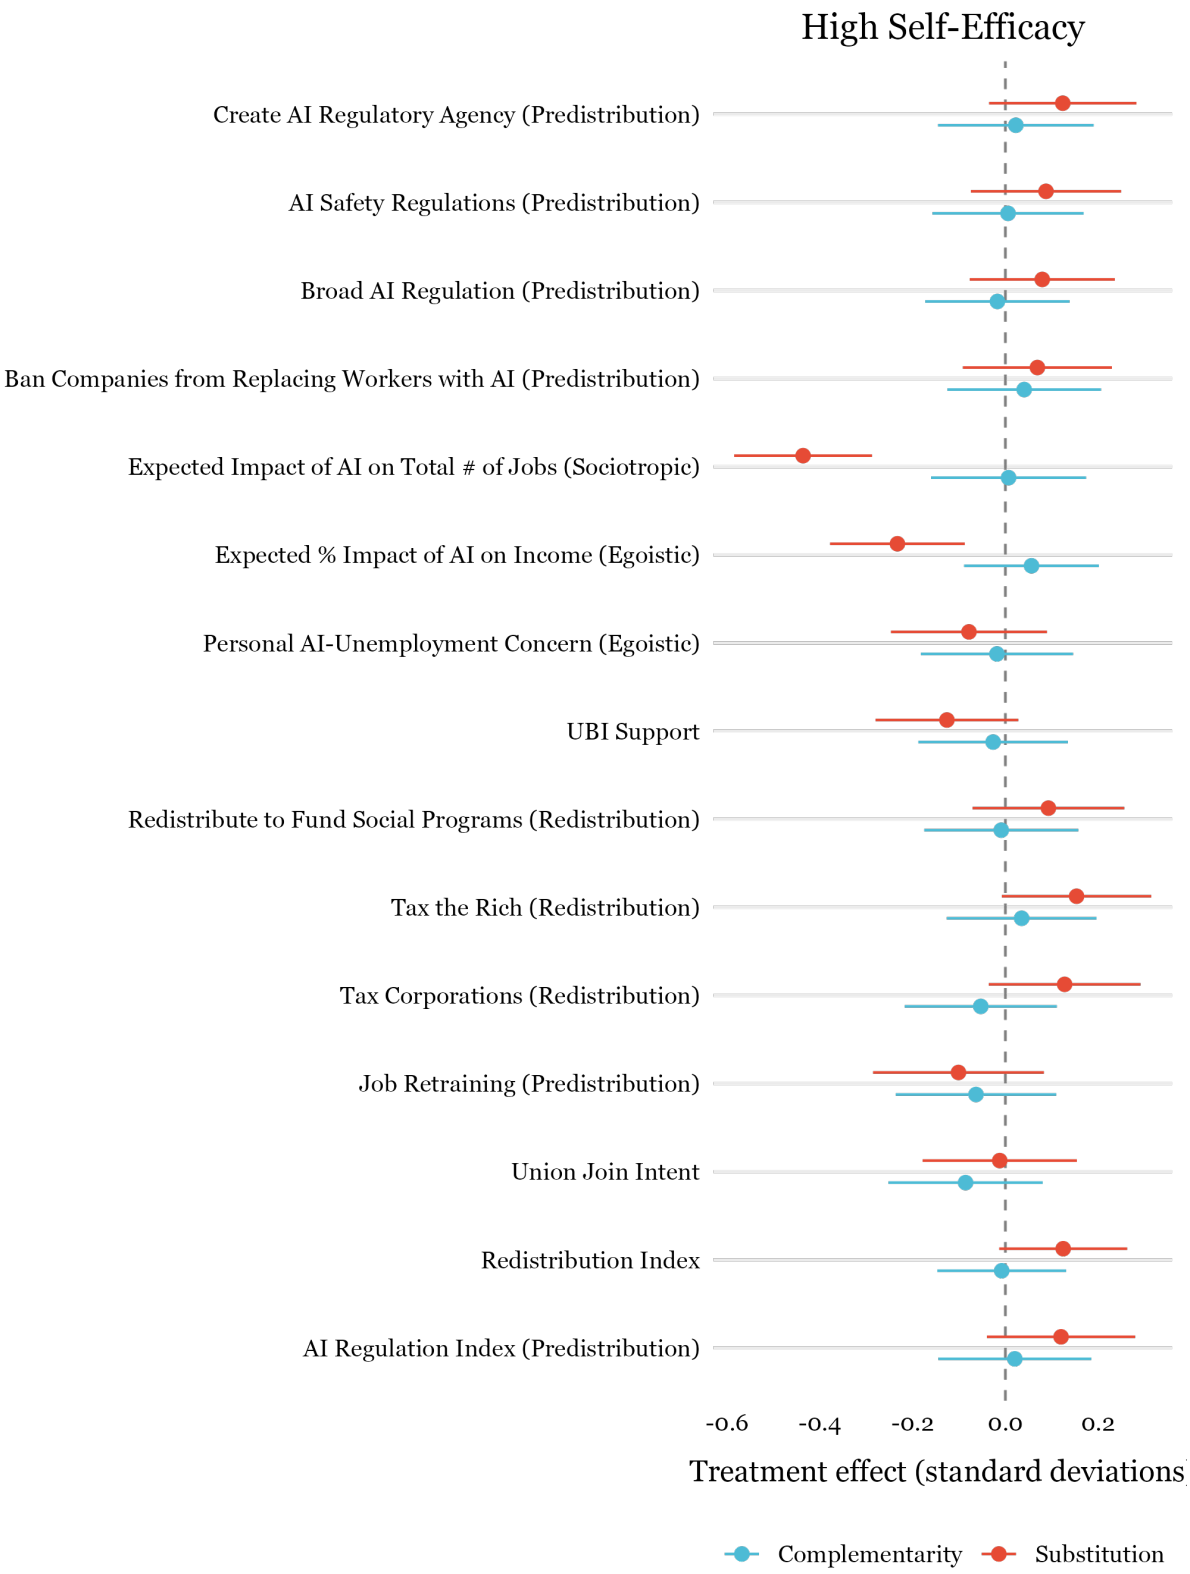


Table C.41: High AI Self-Efficacy: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.125+	−0.126	−0.101	0.120	0.153+	0.128	−0.012
	(0.070)	(0.078)	(0.094)	(0.081)	(0.082)	(0.083)	(0.085)
Complementarity	−0.008	−0.026	−0.063	0.020	0.035	−0.053	−0.086
	(0.071)	(0.082)	(0.088)	(0.084)	(0.082)	(0.084)	(0.085)
Num.Obs.	861	850	859	859	860	861	836
RMSE	0.85	0.95	1.11	0.99	0.98	1.00	1.00

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.42: High AI Self-Efficacy: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.080	0.069	0.088	0.124
	(0.080)	(0.082)	(0.083)	(0.081)
Complementarity	−0.017	0.041	0.006	0.023
	(0.079)	(0.084)	(0.083)	(0.086)
Num.Obs.	852	859	859	858
RMSE	0.94	1.00	0.99	1.00

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.43: High AI Self-Efficacy: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−8.561**	−0.436***	−0.079
	(2.722)	(0.076)	(0.086)
Complementarity	2.069	0.007	−0.018
	(2.723)	(0.085)	(0.084)
Num.Obs.	861	861	862
RMSE	32.23	0.96	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.8 Social Mobility

Downward Mobility

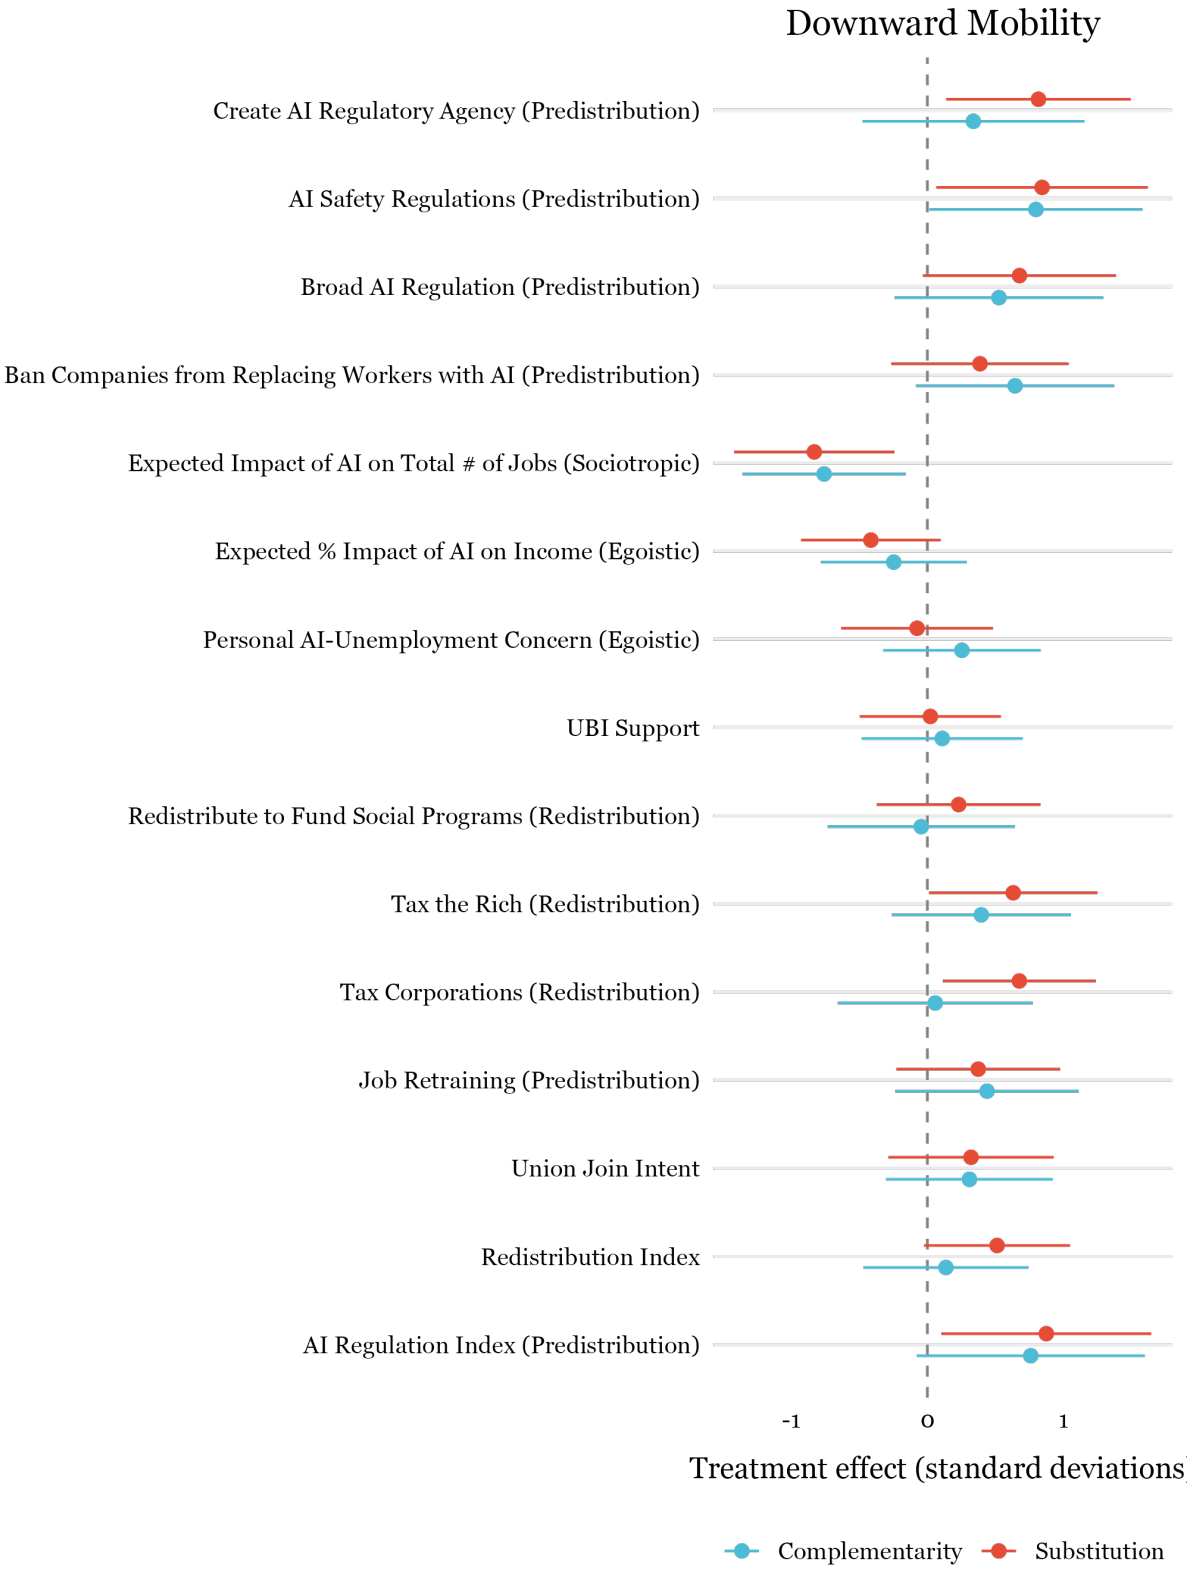


Table C.44: Downward Mobility: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.511+	0.022	0.373	0.872*	0.630*	0.674*	0.320
	(0.269)	(0.260)	(0.302)	(0.386)	(0.310)	(0.282)	(0.304)
Complementarity	0.136	0.109	0.438	0.759+	0.396	0.058	0.308
	(0.304)	(0.297)	(0.338)	(0.420)	(0.330)	(0.359)	(0.307)
Num.Obs.	76	76	76	76	76	76	74
RMSE	0.81	0.87	1.00	1.12	0.91	0.95	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.45: Downward Mobility: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.676+	0.386	0.841*	0.815*
	(0.356)	(0.327)	(0.390)	(0.340)
Complementarity	0.525	0.643+	0.796*	0.337
	(0.384)	(0.365)	(0.393)	(0.409)
Num.Obs.	76	76	76	76
RMSE	1.04	1.10	1.07	1.11

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.46: Downward Mobility: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−15.238	−0.830**	−0.076
	(9.449)	(0.295)	(0.280)
Complementarity	−9.083	−0.758*	0.253
	(9.871)	(0.301)	(0.290)
Num.Obs.	75	76	76
RMSE	31.15	0.89	1.00

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Stable Mobility

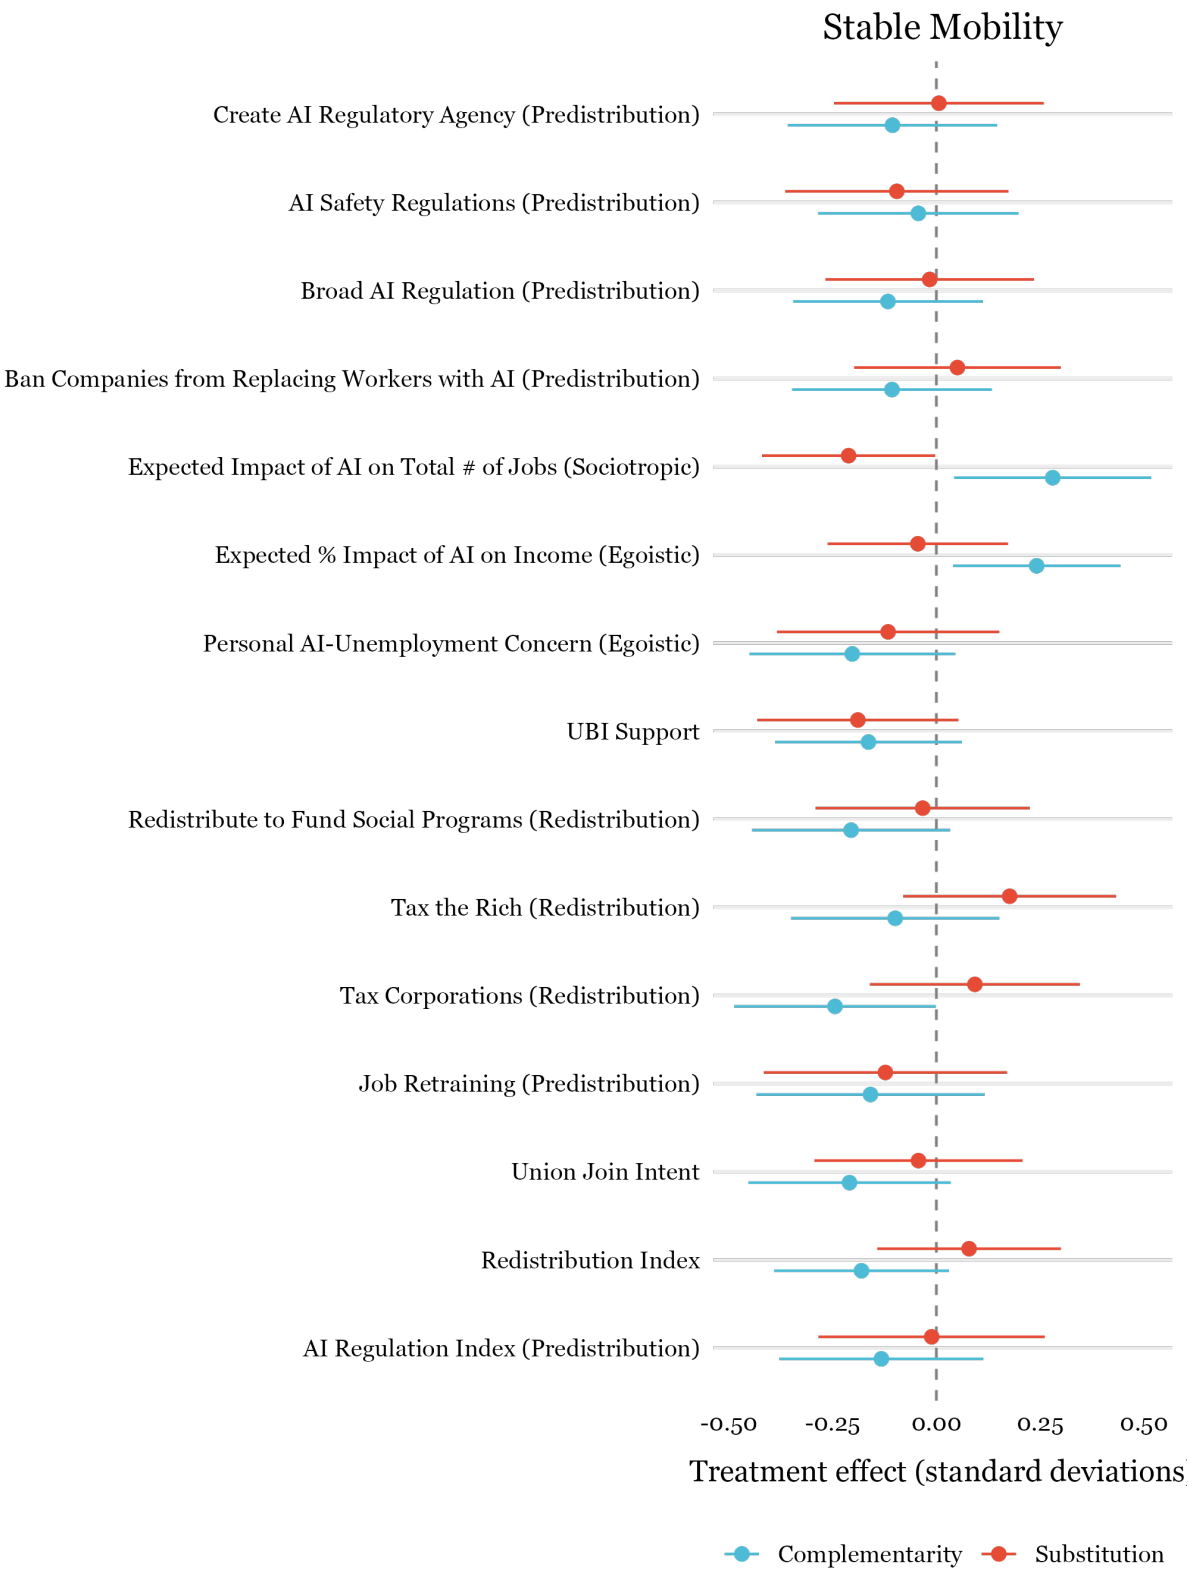


Table C.47: Stable Mobility: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.079	−0.189	−0.123	−0.011	0.177	0.093	−0.043
	(0.113)	(0.123)	(0.149)	(0.139)	(0.130)	(0.129)	(0.127)
Complementarity	−0.180+	−0.164	−0.159	−0.132	−0.099	−0.244*	−0.209+
	(0.107)	(0.114)	(0.140)	(0.125)	(0.128)	(0.124)	(0.124)
Num.Obs.	379	376	378	378	378	379	368
RMSE	0.88	0.94	1.18	1.06	1.03	1.00	0.99

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.48: Stable Mobility: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	−0.016	0.051	−0.095	0.006
	(0.128)	(0.127)	(0.137)	(0.129)
Complementarity	−0.117	−0.107	−0.043	−0.106
	(0.116)	(0.122)	(0.123)	(0.128)
Num.Obs.	373	378	378	377
RMSE	0.96	1.00	1.04	1.03

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.49: Stable Mobility: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−1.642	−0.211*	−0.116
	(4.061)	(0.106)	(0.136)
Complementarity	8.881*	0.280*	−0.203
	(3.768)	(0.121)	(0.126)
Num.Obs.	380	379	380
RMSE	30.18	0.92	1.05

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Upward Mobility

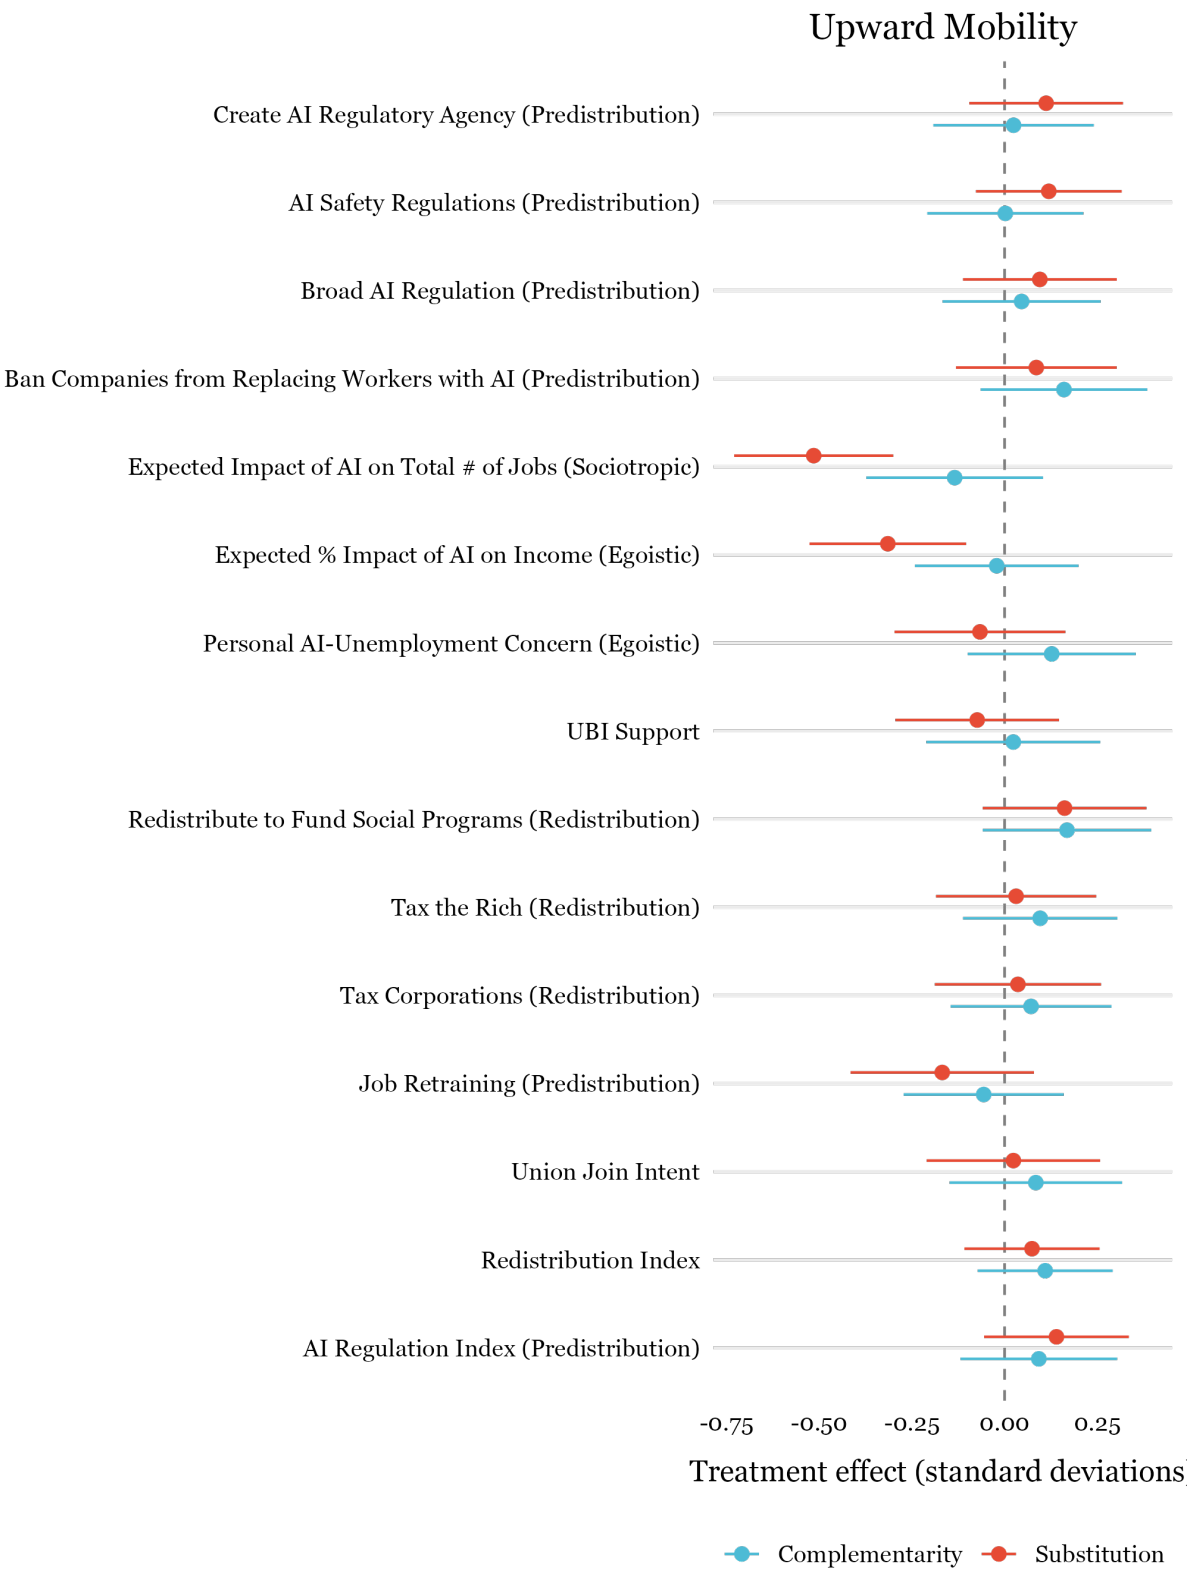


Table C.50: Upward Mobility: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.074	−0.074	−0.168	0.140	0.031	0.036	0.024
	(0.093)	(0.112)	(0.126)	(0.099)	(0.110)	(0.114)	(0.119)
Complementarity	0.110	0.023	−0.056	0.092	0.096	0.071	0.084
	(0.093)	(0.119)	(0.110)	(0.108)	(0.106)	(0.110)	(0.118)
Num.Obs.	463	454	462	462	463	462	445
RMSE	0.82	0.98	1.05	0.91	0.95	0.99	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.51: Upward Mobility: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.095	0.086	0.119	0.112
	(0.105)	(0.110)	(0.100)	(0.105)
Complementarity	0.046	0.160	0.002	0.024
	(0.108)	(0.114)	(0.107)	(0.110)
Num.Obs.	457	462	462	462
RMSE	0.92	0.97	0.92	0.94

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.52: Upward Mobility: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−11.574**	−0.515***	−0.066
	(3.950)	(0.109)	(0.117)
Complementarity	−0.782	−0.135	0.127
	(4.127)	(0.121)	(0.115)
Num.Obs.	463	463	463
RMSE	35.17	0.99	1.03

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

C.9 AI Risk

Low AI Occupational Risk

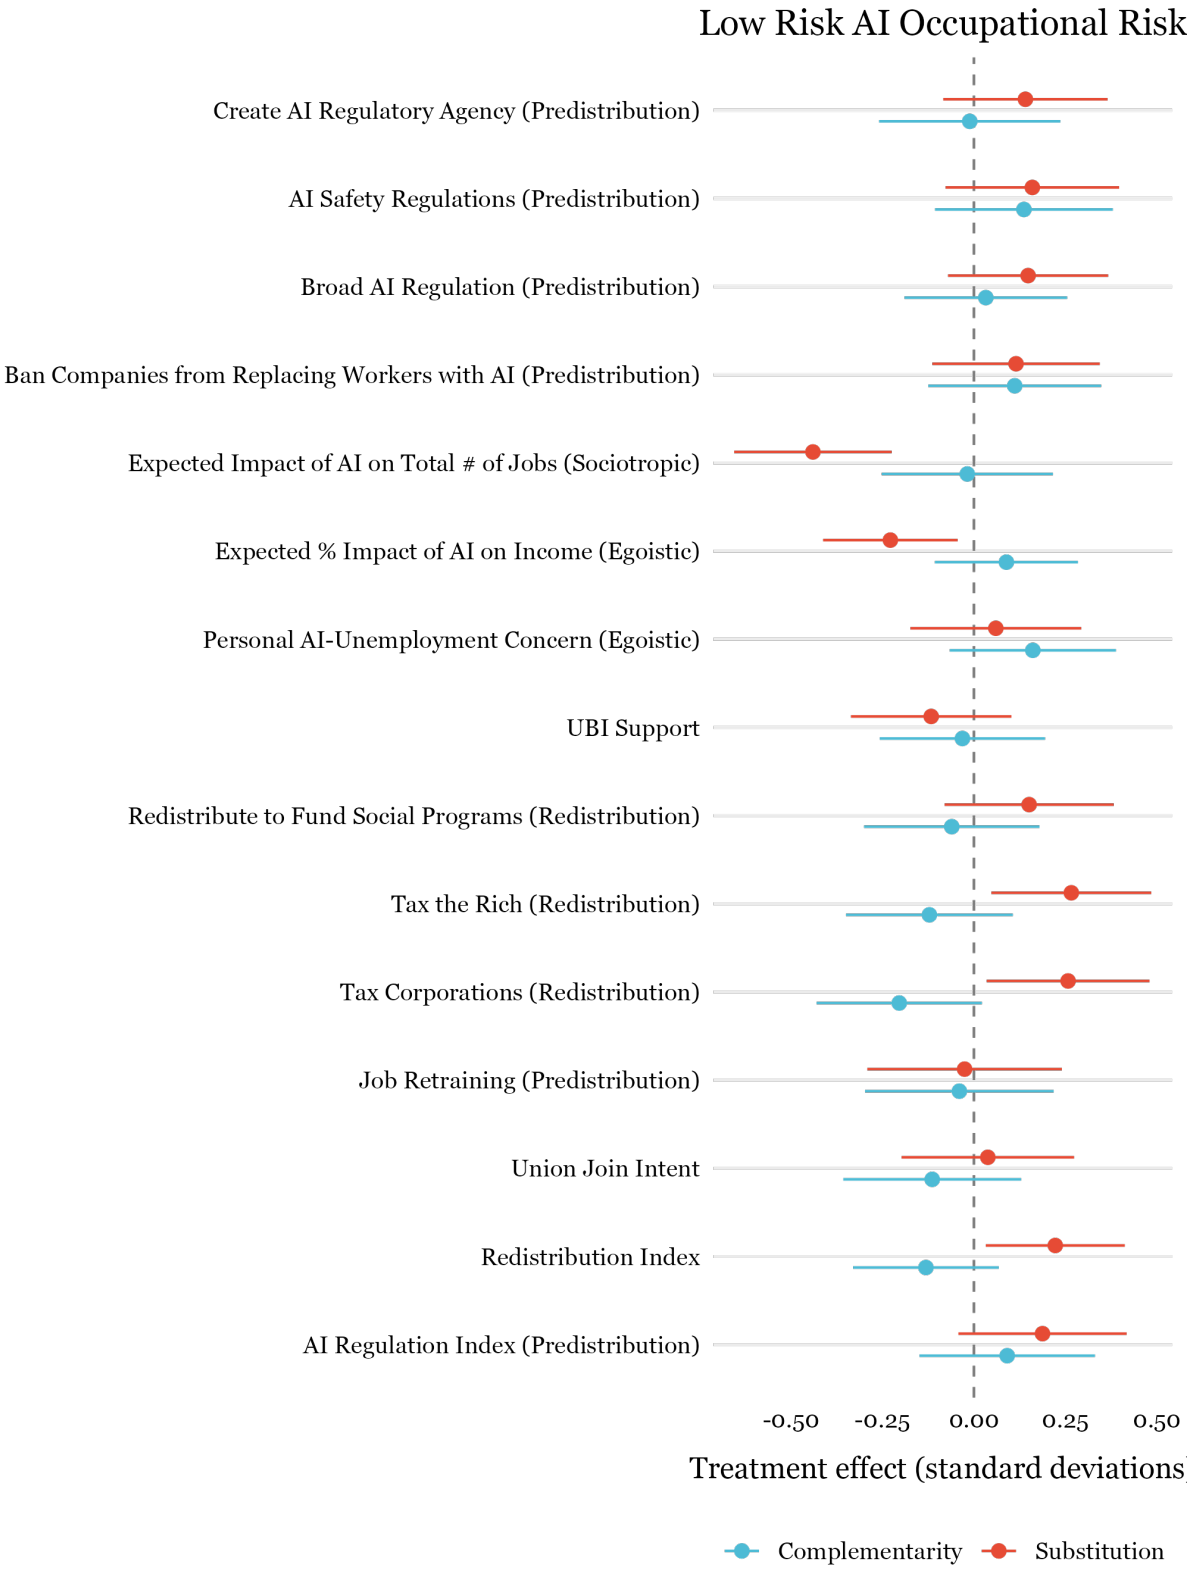


Table C.53: Low AI Occupational Risk: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.222*	−0.117	−0.025	0.187	0.266*	0.257*	0.038
	(0.097)	(0.112)	(0.135)	(0.117)	(0.111)	(0.113)	(0.120)
Complementarity	−0.131	−0.031	−0.040	0.091	−0.121	−0.204+	−0.114
	(0.101)	(0.115)	(0.131)	(0.122)	(0.116)	(0.115)	(0.124)
Num.Obs.	444	438	443	442	444	443	432
RMSE	0.85	0.94	1.16	1.01	0.97	0.98	1.03

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.54: Low AI Occupational Risk: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.148	0.115	0.160	0.141
	(0.111)	(0.116)	(0.121)	(0.114)
Complementarity	0.032	0.111	0.137	−0.011
	(0.113)	(0.120)	(0.124)	(0.126)
Num.Obs.	438	442	442	441
RMSE	0.93	1.01	1.02	1.03

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.55: Low AI Occupational Risk: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−8.388*	−0.440***	0.060
	(3.438)	(0.110)	(0.119)
Complementarity	3.263	−0.018	0.161
	(3.655)	(0.119)	(0.116)
Num.Obs.	444	445	445
RMSE	30.18	0.97	1.02

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

High AI Occupational Risk

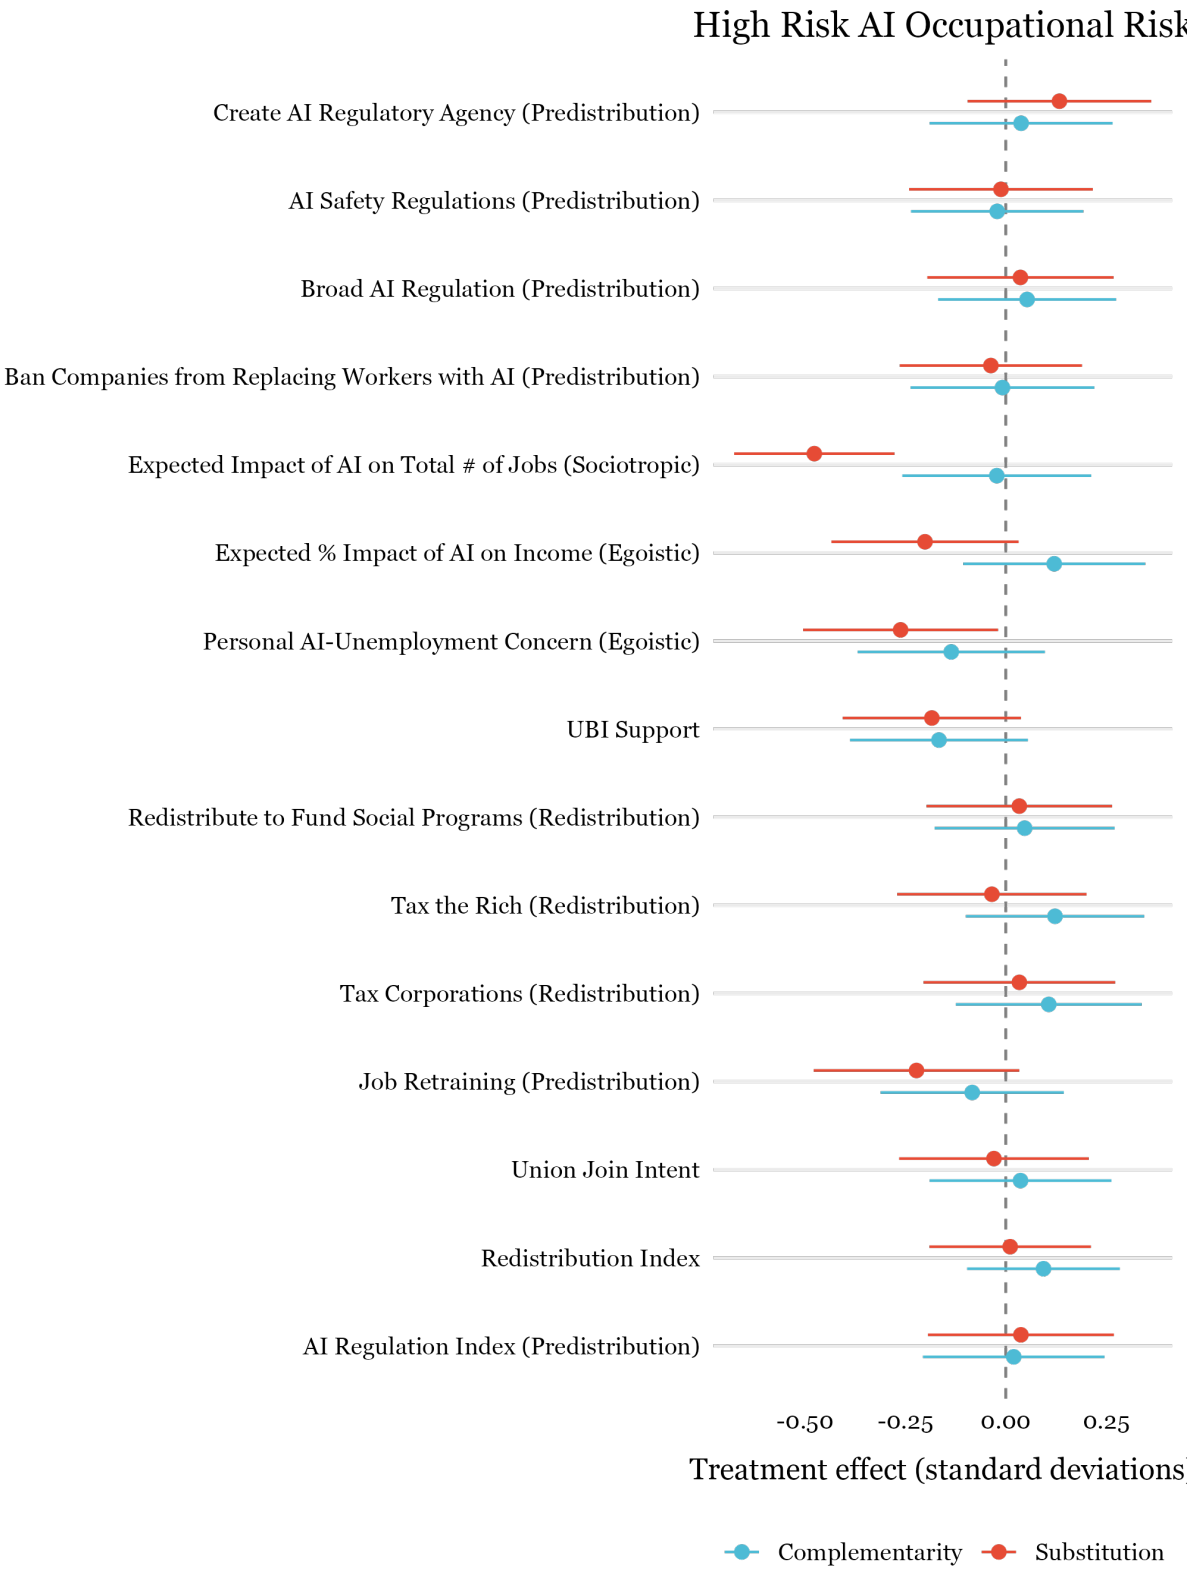


Table C.56: High AI Occupational Risk: Additional Outcomes (Standardized)

	Redistribution Index	UBI	Job Retraining	Regulation Index	Tax the Rich	Tax Corporations	Union Interest
Substitution	0.011	−0.184	−0.222+	0.038	−0.035	0.034	−0.029
	(0.102)	(0.113)	(0.130)	(0.118)	(0.120)	(0.121)	(0.120)
Complementarity	0.094	−0.166	−0.084	0.020	0.123	0.107	0.037
	(0.097)	(0.113)	(0.116)	(0.115)	(0.113)	(0.118)	(0.115)
Num.Obs.	435	429	434	435	434	435	417
RMSE	0.85	0.95	1.08	1.00	0.99	1.02	0.99

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.57: High AI Occupational Risk: Regulation Sub-Items (Standardized)

	General AI Regulation	AI Layoff Limits	Safety Regulation	AI Agency Creation
Substitution	0.037	−0.037	−0.012	0.134
	(0.118)	(0.115)	(0.116)	(0.116)
Complementarity	0.053	−0.008	−0.021	0.038
	(0.113)	(0.116)	(0.109)	(0.116)
Num.Obs.	430	435	435	435
RMSE	0.98	0.99	0.97	0.99

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table C.58: High AI Occupational Risk: Mediators

	Expected \% Income Impact (EFI)	Expected Impact on Jobs	AI Unemployment Concern
Substitution	−7.389+	−0.477***	−0.262*
	(4.358)	(0.102)	(0.123)
Complementarity	4.436	−0.022	−0.136
	(4.238)	(0.120)	(0.119)
Num.Obs.	435	434	435
RMSE	35.46	0.93	1.04

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

D Robustness Checks: Results by Party

Republicans respond to both treatments with a desire to reduce job retraining programs for those who lose their jobs to AI. To investigate whether this result is a manifestation of statistical noise, I run a number of interaction and bootstrap models in both simple form (dichotomous) and saturated (multivariate with all controls) form.

Table D.1: Partisan Support for Job Retraining, Simple Model

	Dependent variable:
	job_retraining_z
pIndependent/Other	−0.165
	p = 0.332
pRepublican	−0.293
	p = 0.031
treatment_groupComplementarity	0.094
	p = 0.475
treatment_groupSubstitution	0.126
	p = 0.342
education_levelNo Degree	0.016
	p = 0.859
pIndependent/Other:treatment_groupComplementarity	−0.182
	p = 0.442
pRepublican:treatment_groupComplementarity	−0.410
	p = 0.035
pIndependent/Other:treatment_groupSubstitution	−0.135
	p = 0.573
pRepublican:treatment_groupSubstitution	−0.517
	p = 0.008
Constant	0.137
	p = 0.144
Observations	953
R^2	0.067
Adjusted R^2	0.058
Residual Std. Error	1.090 (df = 943)
F Statistic	7.554^{***} (df = 9; 943)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table D.2: Saturated Interaction Model

	Dependent variable:
	job_retraining_z
pIndependent/Other	−0.225 p = 0.224
pRepublican	−0.388 p = 0.007
treatment_groupComplementarity	−0.004 p = 0.980
treatment_groupSubstitution	0.004 p = 0.980
education_levelNo Degree	0.017 p = 0.865
prolific_ethnicityEast Asian	0.149 p = 0.578
prolific_ethnicityLatino/Hispanic	0.197 p = 0.241
prolific_ethnicityMixed	0.608 p = 0.139
prolific_ethnicityNative American or Alaskan Native	0.821 p = 0.078
prolific_ethnicityOther (please feel free to let us know your ethnicity via email)	−0.483 p = 0.302
prolific_ethnicitySouth Asian	−0.080 p = 0.798
prolific_ethnicitySouth East Asian	0.395 p = 0.154
prolific_ethnicityWhite	0.009 p = 0.949
prolific_age	0.007 p = 0.167
prolific_sexMale	−0.123 p = 0.106
log(income_corr + 1)	−0.090 p = 0.071
risk	0.550 p = 0.042
pIndependent/Other:treatment_groupComplementarity	−0.032 p = 0.900
pRepublican:treatment_groupComplementarity	−0.283 p = 0.170
pIndependent/Other:treatment_groupSubstitution	−0.108

Table D.3: Simple Repub Job Retraining

	term	estimate	boot_se	bca_lo	bca_hi
pRepublican:Substitution	Republican x Substitution	-0.515	0.183	-0.867	-0.153

Table D.4: Saturated Rep Job Retraining

	term	estimate	boot_se	bca_lo	bca_hi
pRepublican:Substitution	Republican x Substitution	-0.380	0.211	-0.790	0.060

Table D.5:

	term	estimate	boot_se	bca_lo	bca_hi
1	Republican x Complementarity	-0.416	0.192	-0.797	-0.026

Table D.6: Saturated Repub Job Retraining 2

	term	estimate	boot_se	bca_lo	bca_hi
1	Republican x Complementarity	-0.283	0.201	-0.665	0.161

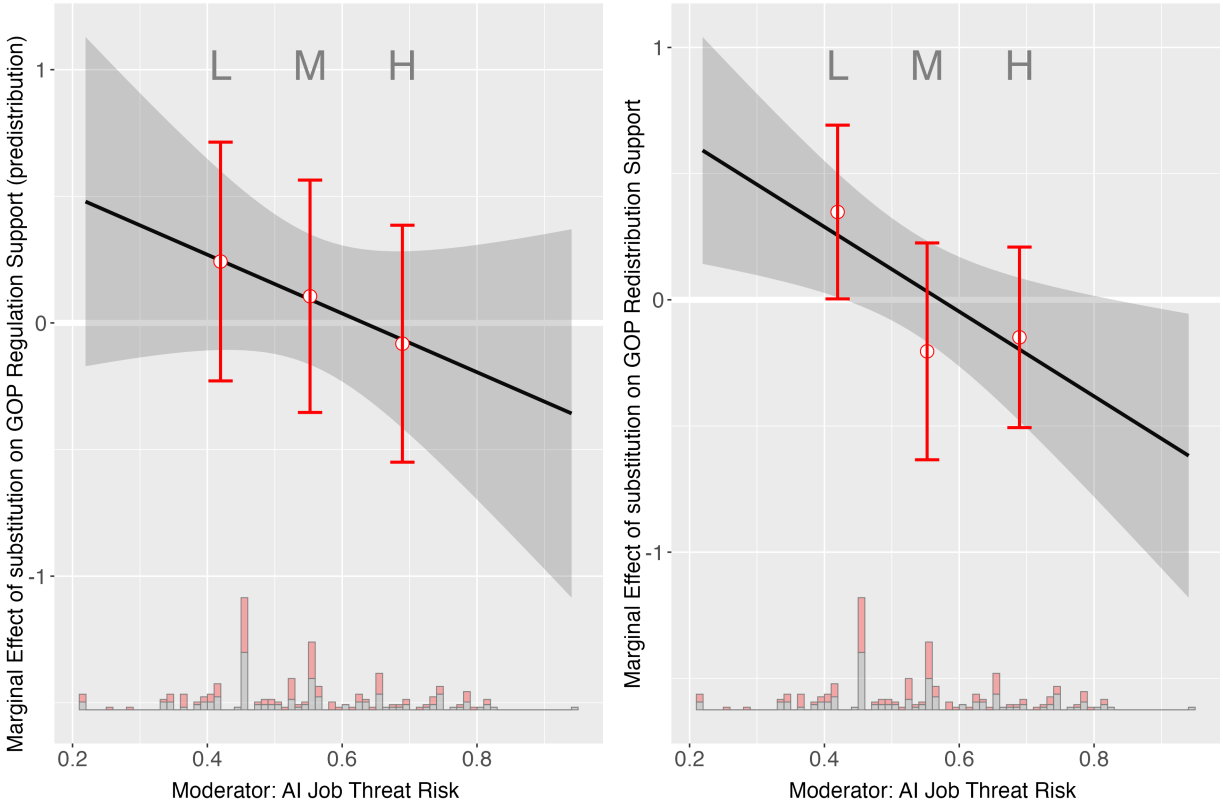


Figure 1: Interflex: Job Risk and Main DVs among Republicans

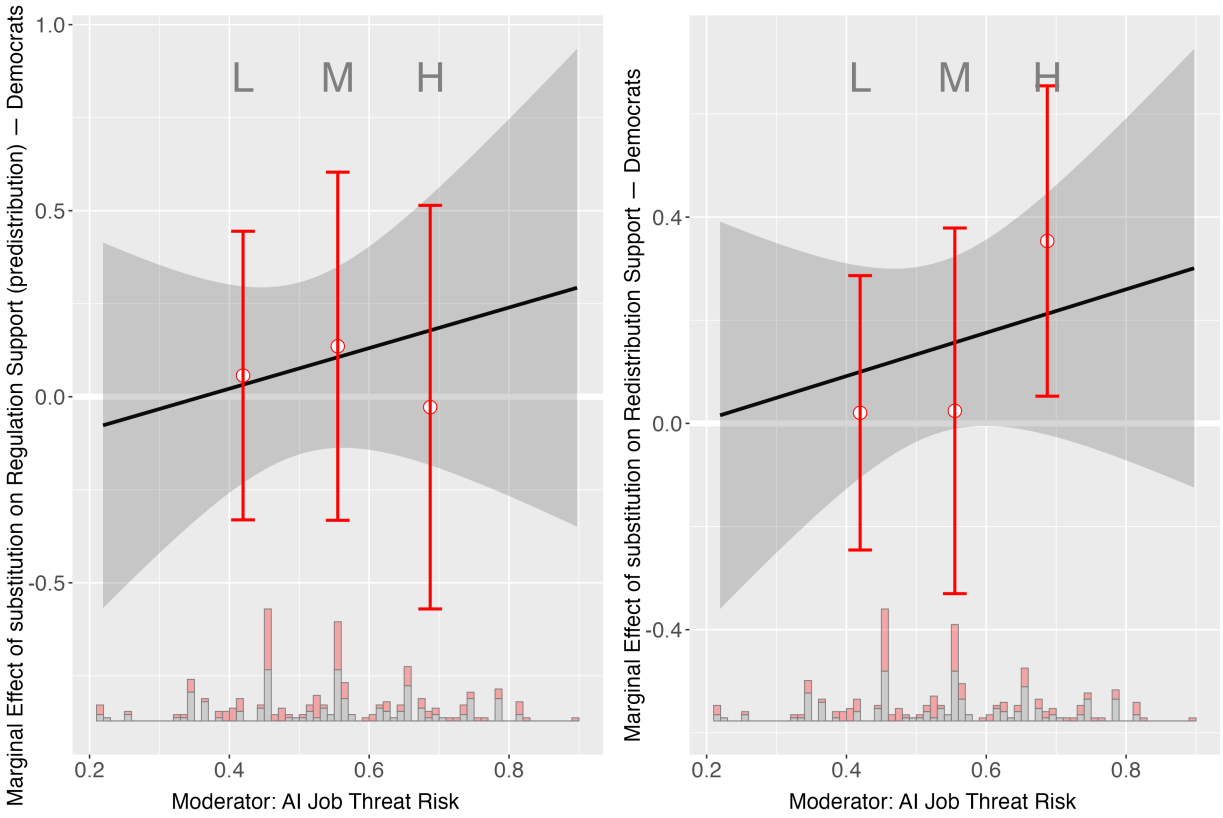


Figure 2: Interflex: Job Risk and Main DVs among Democrats

E Robustness Checks: Results by Education Level

T-test results show that when my sample is subset to those with no degree, individuals shown the AI-threat prime prefer to increase AI regulation, which is a key form of predistribution. However, this portion of my sample is comparatively small (n = 127), and additional checks are required to identify if they are meaningful. To test whether these results are robust, I first compare the demographic makeup of the uneducated to the educated in my sample to show that they are not over-representing certain groups groups. Then, I show the results of interaction models and bootstrap models.

My sample is heavily weighted toward individuals with a college education, as 69% hold a BA or higher. As such, I conduct a number of robustness tests to verify the strength of this result. I first test whether this result holds if Associates Degree holders are considered to *not* have a degree and find that my results differ depending on how associate's and technical degrees are defined. In the US, associate's and technical degrees typically require two years of post-secondary schooling at either a community college for a vocational school. Significant policy preference shifts are found in when AA holders are considered degree holders (see figure 8), but not when AA holders are considered non-degree holders. When AA holders are simply excluded from a dichotomized analysis, significant policy preference where non-degree holders prefer regulation in the face of AI-threat frames re-emerge.

I argue that associate's and technical degree holders more conceptually similar to BA holders than they are to non-degree holders and ought to be considered degree holders, though it is worth noting that associate's degree holders do not neatly fit into a dichotomized view of educational attainment. They are *dissimilar* to BA holders in the sense that their education is often vocational in nature and ii) they do not have a traditional "college experience." In fact, they are unique from both BA holders *and* high school graduates in the sense that they are far more likely than any group to be political independents (27%) and have far and away the strongest predisposition toward AI regulation in the sample.

However, Associate's and technical degree holders are *similar* to BA holders in that their i) their education allows them to command higher salaries than non-degree holders (7.2% higher in this sample), ii) they are often qualified to perform a job that high school graduates are not, with the most common occupation of AA holders in my sample being Managers as compared with the most common occupation of non-holders being Customer Service Representative, and iii) they obtained a degree at a post-secondary institution. For the purposes of this study, I consider AA or technical degree holders to be degree holders.

Table E.1: Average Sample Characteristics by Education Level

	No Degree	Degree
Age	41	39
AI Risk Score	0.57	0.54
Monthly Income	$3367	$4429
Sex (% Female)	54%	53%
Democrat	38%	45%
Independent/Other	24%	18%
Republican	38%	37%

Table E.2: Ethnicity by Education Level

Ethnicity	Degree	No Degree
Black/African American	58 (7.6%)	19 (10.2%)
East Asian	22 (2.9%)	0 (0.0%)
Latino/Hispanic	95 (12.4%)	25 (13.4%)
Mixed	6 (0.8%)	3 (1.6%)
Native American or Alaskan Native	5 (0.7%)	1 (0.5%)
Other	5 (0.7%)	1 (0.5%)
South Asian	14 (1.8%)	1 (0.5%)
South East Asian	18 (2.4%)	2 (1.1%)
White	502 (65.7%)	122 (65.6%)
NA	39 (5.1%)	12 (6.5%)

Additionally, degree holders support slightly more AI regulation than non degree holders on average, with most of this being driven by the fact that AA and technical degree holders are by far the most interested in regulating AI in the absence of messaging.

Next, I compare the relationship between education and AI regulation attitudes in my own sample's control group to external Pew data. I find that increased education predicts increased support for AI regulation in both my sample and Pew's American Trends Panels (ATP), suggesting my sample is an accurate portrayal of the workforce's attitudes toward AI, though it is worth noting that AA holders do not have

Table E.3: Mean Regulation Index by Education Level (Control Only)

Education Level	Mean
High school diploma or GED	0.64
Some college, but no degree	0.73
Associates or technical degree	0.97
Bachelor's degree	0.7
Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, DDS etc.)	0.73

Table E.4: Partisanship by Education Level (%)

Education Level	Republican	Democrat	Independent
High school diploma or GED	39.7	39.7	20.7
Some college, but no degree	39.5	39.5	21.0
Associates or technical degree	40.8	32.0	27.2
Bachelor's degree	34.9	49.7	15.4
Graduate or professional degree	42.0	44.3	13.7

Table E.5: Mean Regulation Index by Education Level (Control Only)

Education Level	Mean
High school diploma or GED	0.64
Some college, but no degree	0.73
Associates or technical degree	0.97
Bachelor's degree	0.7
Graduate or professional degree (MA, MS, MBA, PhD, JD, MD, DDS etc.)	0.73

the highest level of baseline AI regulation support in Pew's sample. Pew frequently conducts American Trends Panels on various issues. In August 2024, they conducted their most recent ATP on attitudes toward Artificial Intelligence. Pew does not directly ask if individuals want more AI regulation, however, but they do so indirectly in asking "Thinking about the use of artificial intelligence (AI) in the United States, are you more concerned that the U.S. government will..." with the answers available being "Go too far regulating its use", "Not go far enough regulating its use", and "Not sure". I consider "Not sure" a neutral answer, "Go to far" as meaning a person wants less regulation (coded as a a value of -1) and "Not go far enough" as meaning they want more AI regulation (coded as +1).

There are some limitations to this comparison. The first is that Pew does not collect occupational code data in ATP polls, so their sample is representative of the entire US, while my sample is solely comprised of the white-collar workers because of their unique vulnerability to AI. Second, their survey was conducted in 2024, prior to the aforementioned partisan realignment with regard to AI, which may impact support for AI regulation in surprising ways. I consider the general similarity of my sample's relationship with AI regulation and Pew's, however, to be indicative of robustness.

Table E.6: Mean AI Regulation Support by Education Level (Pew ATP August 2024)

Education Level	Regulation Support	n
Less than HS	.07	213
HS	.30	1106
Some college	.40	1029
Associate	.45	486
Bachelor	.55	1240
Postgrad	.51	931

Table E.7: Pew ATP Data: Education and AI Regulation Attitudes (OLS; education ordinal 1-6)

	Dependent variable:
	Support for AI Regulation
edu_ord	0.052*** (0.008)
age_ord	0.044*** (0.011)
genderWoman	−0.044** (0.022)
genderOther	0.263* (0.143)
race_ethBlack non-Hisp	−0.262*** (0.037)
race_ethHispanic	−0.112*** (0.032)
race_ethOther	−0.135** (0.058)
race_ethAsian non-Hisp	−0.185*** (0.038)
income_ord	0.015*** (0.004)
partyRep	−0.126*** (0.030)
partyDem	0.134*** (0.029)
partyOther	−0.024 (0.037)
Constant	0.117** (0.052)
N	5,005
R^2	0.050
Adj. R^2	0.047
Res. Std. Err.	0.773 (df = 4992)

Note: *p<0.1; **p<0.05; ***p<0.01

Table E.8: Present Study: Baseline AI Regulation Attitudes

	Dependent variable:
	AI Regulation Index
Education	−0.030 (0.050)
Age	−0.008 (0.007)
Male	−0.093 (0.115)
East Asian	−0.398 (0.391)
Latino/Hispanic	0.325 (0.241)
Mixed	−0.078 (0.714)
Native American or Alaskan Native	0.608 (0.994)
Other	0.422 (0.527)
South Asian	0.597 (0.442)
South East Asian	0.294 (0.392)
White	−0.013 (0.199)
log(income)	−0.111 (0.081)
Independent/Other	−0.307* (0.161)
Republican	−0.481*** (0.126)
Constant	1.569** (0.796)
N	302
R^2	0.097
Adj. R^2	0.053

Note: *p<0.1; **p<0.05; ***p<0.01

Given the relatively small amount of non degree holders in my sample (n =126), I run interaction and bootstrap models testing the extent to which education moderates the relationship between between AI-threat frames and support for regulation. I find that in simple interaction models, the interaction between education and the AI-threat treatment is not significant at conventional levels (p = .108), but if the baseline is shifted from the control treatment to the Complementarity treatment, the relevant p value drops to .004. Similar dynamics are observed in saturated models, wherein simple models show no statistically significant interaction between education and the substitution treatment but the relevant p value drops to .06 when the baseline is moved to the complementarity treatment. Similarly, simple bootstrap models show a statistically significant interaction but saturated bootstrap models are slightly below conventional significance levels.

Table E.9: Pew Data: Education and AI Regulation Attitudes (OLS; education ordinal 1-6)

	Dependent variable:
	Support for AI Regulation
edu_ord	0.052^{***}
	(0.008)
age_ord	0.044^{***}
	(0.011)
genderWoman	-0.044^{**}
	(0.022)
genderOther	0.263^{*}
	(0.143)
race_ethBlack non-Hisp	-0.262^{***}
	(0.037)
race_ethHispanic	-0.112^{***}
	(0.032)
race_ethOther	-0.135^{**}
	(0.058)
race_ethAsian non-Hisp	-0.185^{***}
	(0.038)
income_ord	0.015^{***}
	(0.004)
partyRep	-0.126^{***}
	(0.030)
partyDem	0.134^{***}
	(0.029)
partyOther	-0.024
	(0.037)
Constant	0.117^{**}
	(0.052)
N	5005
Observations	5,005
R^2	0.050
Adjusted R^2	0.047
Residual Std. Error	0.773 (df = 4992)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table E.10: Present Study: Baseline AI Regulation Attitudes

	Dependent variable:
	AI Regulation Index
education	−0.030
	(0.050)
prolific_age	−0.008
	(0.007)
prolific_sexMale	−0.093
	(0.115)
prolific_ethnicityEast Asian	−0.398
	(0.391)
prolific_ethnicityLatino/Hispanic	0.325
	(0.241)
prolific_ethnicityMixed	−0.078
	(0.714)
prolific_ethnicityNative American or Alaskan Native	0.608
	(0.994)
prolific_ethnicityOther	0.422
	(0.527)
prolific_ethnicitySouth Asian	0.597
	(0.442)
prolific_ethnicitySouth East Asian	0.294
	(0.392)
prolific_ethnicityWhite	−0.013
	(0.199)
log(income_corr + 1)	−0.111
	(0.081)
pIndependent/Other	−0.307*
	(0.161)
pRepublican	−0.481***
	(0.126)
Constant	1.569**
	(0.796)
Observations	302
R^2	0.097
Adjusted R^2	0.053
Note:	*p<0.1; **p<0.05; ***p<0.01

Table E.11: Simple Interaction Model

	Dependent variable:
	regulation_index_z
education_levelNo Degree	−0.038
	p = 0.790
treatment_groupComplementarity	0.087
	p = 0.333
treatment_groupSubstitution	0.056
	p = 0.534
education_levelNo Degree:treatment_groupComplementarity	−0.275
	p = 0.180
education_levelNo Degree:treatment_groupSubstitution	0.326
	p = 0.108
Constant	0.007
	p = 0.907
Observations	953
R^2	0.012
Adjusted R^2	0.007
Residual Std. Error	1.019 (df = 947)
F Statistic	2.249** (df = 5; 947)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table E.12: Simple Interaction Model, Complementarity baseline

	Dependent variable:
	regulation_index_z
education_levelNo Degree	−0.314
	p = 0.032
treatment_groupControl	−0.087
	p = 0.333
treatment_groupSubstitution	−0.031
	p = 0.733
education_levelNo Degree:treatment_groupControl	0.275
	p = 0.180
education_levelNo Degree:treatment_groupSubstitution	0.601
	p = 0.004
Constant	0.095
	p = 0.138
Observations	953
R^2	0.012
Adjusted R^2	0.007
Residual Std. Error	1.019 (df = 947)
F Statistic	2.249** (df = 5; 947)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table E.13: Saturated Interaction Model

	Dependent variable:
	regulation_index_z
education_levelNo Degree	0.023 p = 0.882
treatment_groupComplementarity	0.121 p = 0.191
treatment_groupSubstitution	0.088 p = 0.343
prolific_ethnicityEast Asian	−0.085 p = 0.726
prolific_ethnicityLatino/Hispanic	0.337 p = 0.028
prolific_ethnicityMixed	0.382 p = 0.304
prolific_ethnicityNative American or Alaskan Native	0.742 p = 0.079
prolific_ethnicityOther (please feel free to let us know your ethnicity via email)	0.060 p = 0.887
prolific_ethnicitySouth Asian	0.396 p = 0.162
prolific_ethnicitySouth East Asian	0.464 p = 0.065
prolific_ethnicityWhite	0.113 p = 0.364
prolific_age	−0.010 p = 0.025
prolific_sexMale	−0.217 p = 0.002
log(income_corr + 1)	−0.030 p = 0.504
pIndependent/Other	−0.335 p = 0.0005
pRepublican	−0.492 p = 0.000
risk	0.460 p = 0.061
education_levelNo Degree:treatment_groupComplementarity	−0.166 p = 0.444
education_levelNo Degree:treatment_groupSubstitution	0.229 p = 0.278
Constant	0.530

Table E.14: Saturated Interaction Model, Complementarity Baseline

	Dependent variable:
	regulation_index_z
education_levelNo Degree	−0.144
	p = 0.360
treatment_groupControl	−0.121
	p = 0.191
treatment_groupSubstitution	−0.033
	p = 0.726
prolific_ethnicityEast Asian	−0.085
	p = 0.726
prolific_ethnicityLatino/Hispanic	0.337
	p = 0.028
prolific_ethnicityMixed	0.382
	p = 0.304
prolific_ethnicityNative American or Alaskan Native	0.742
	p = 0.079
prolific_ethnicityOther	0.060
	p = 0.887
prolific_ethnicitySouth Asian	0.396
	p = 0.162
prolific_ethnicitySouth East Asian	0.464
	p = 0.065
prolific_ethnicityWhite	0.113
	p = 0.364
prolific_age	−0.010
	p = 0.025
prolific_sexMale	−0.217
	p = 0.002
log(income_corr + 1)	−0.030
	p = 0.504
pIndependent/Other	−0.335
	p = 0.0005
pRepublican	−0.492
	p = 0.000
risk	0.460
	p = 0.061
education_levelNo Degree:treatment_groupControl	0.166
	p = 0.444
education_levelNo Degree:treatment_groupSubstitution	0.395
	p = 0.066
Constant	0.651
	p = 0.156
Observations	844
R^2	0.095
Adjusted R^2	0.074
Residual Std. Error	0.987 (df = 824)
F Statistic	4.558^{***} (df = 19; 824)
Note:	*p<0.1; **p<0.05; ***p<0.01

E.1 Comparing Results by Education with Other Studies

My findings are somewhat aligned with the results of Haslberger et al. (2025). While they do not show that those *without* a degree want more AI regulation, they show that those *with* a degree want *less* regulation in response to using AI. In other words, they show what might be considered the expected inverse effect of what I find. The relevant table etc. is on Figure F2 of their appendix. It is worth noting that they also find that those with no degree are more

concerned about losing their job to an AI-skilled human post-treatment. However, their treatment is different than mine (users *use* ChatGPT in their experiment).

In their supplementary materials, Borwein et al., (2024) report that Republicans primed with AI threats do, like in my sample, support *less universal basic income*. However, Republicans in my study react differently with regard to support for job retraining. In their sample, Republicans faced with an AI shock show increased support for job retraining after taking part in their conjoint experiment which highlights the role of AI in inflating prices (See Borwein et al., (2024), appendix figure G1). The present study, on the other hand, shows that video primes of the AI threat have the opposite effect on Republicans by making them **less supportive of job retraining** for those laid off due to AI.

This result begs the question of whether Republicans in this sample reacted to the AI threat differently in terms of their personal and sociotropic economic prospects.

Table E.15: No Degree Regulation Index Bootstrap; Simple Model

	term	estimate	boot_se	bca_lo	bca_hi
treatment_groupSubstitution	Substitution (vs control)	0.382	0.183	0.029	0.739

Table E.16: Saturated Model

	term	estimate	boot_se	bca_lo	bca_hi
treatment_groupSubstitution	Substitution (vs control)	0.326	0.205	-0.040	0.780

F Average Causal Mediation Effects (ACME)

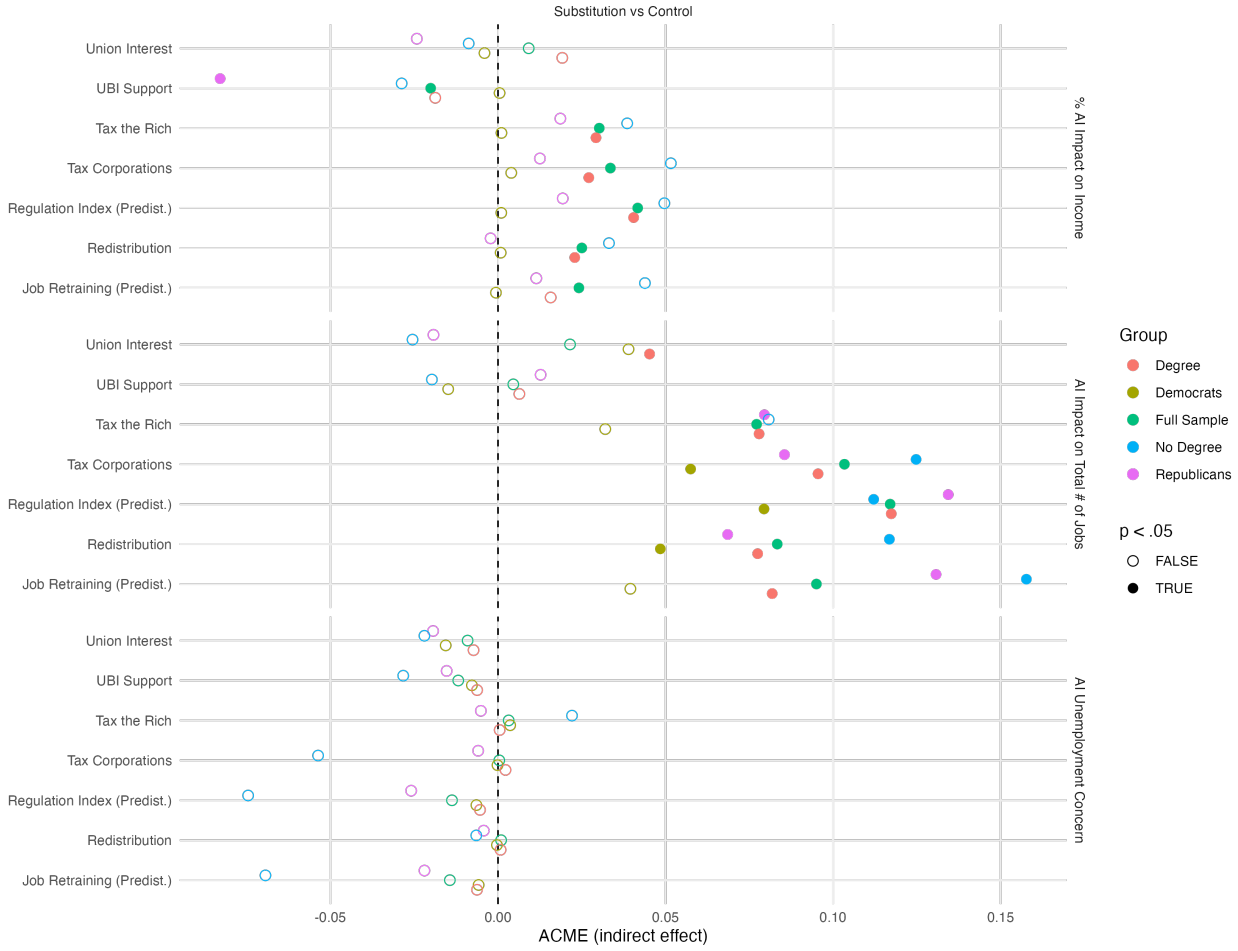


Figure 3: Full Sample, Substitution vs Control

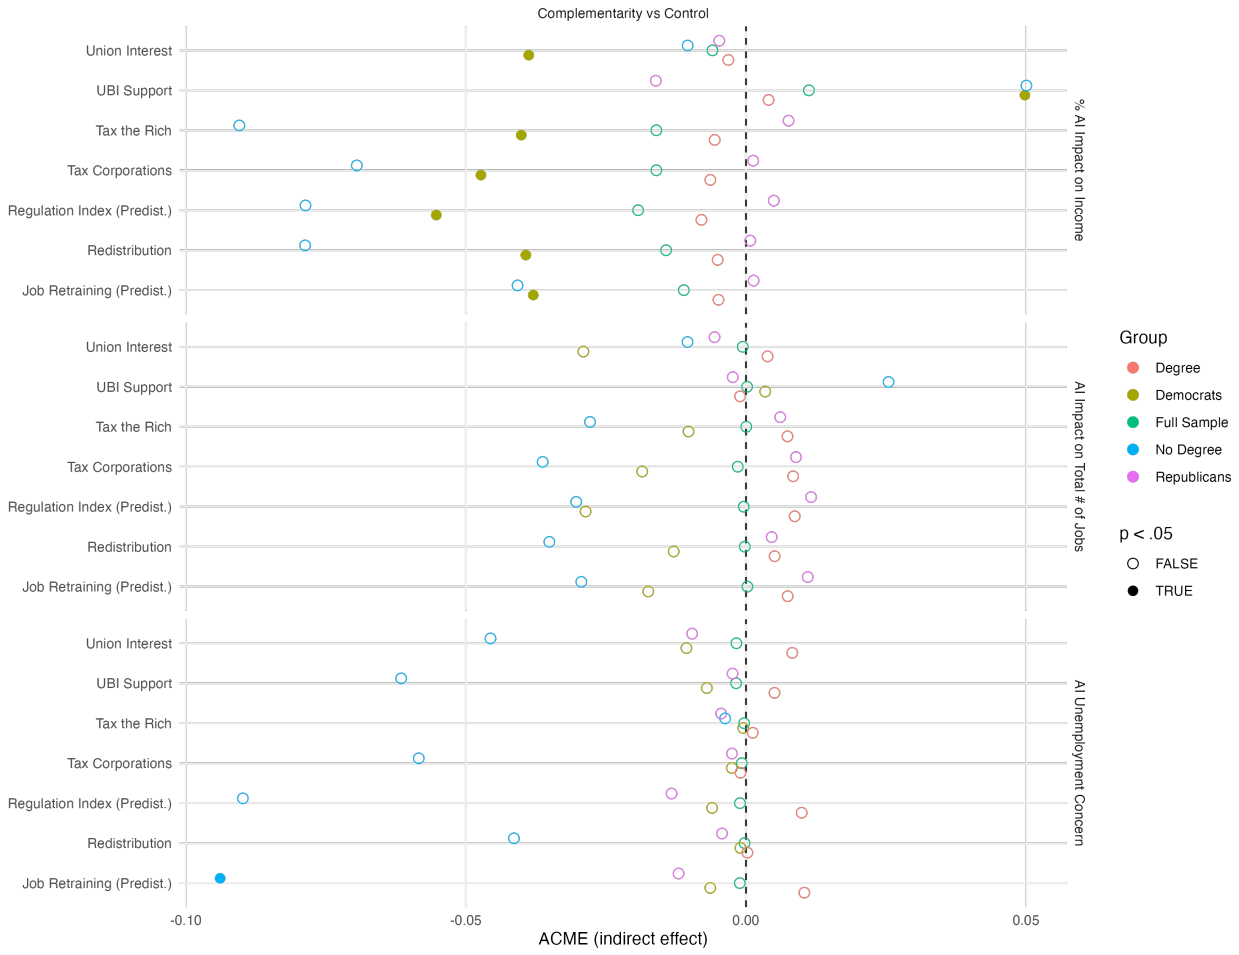


Figure 4: Full Sample, Complementarity vs Control

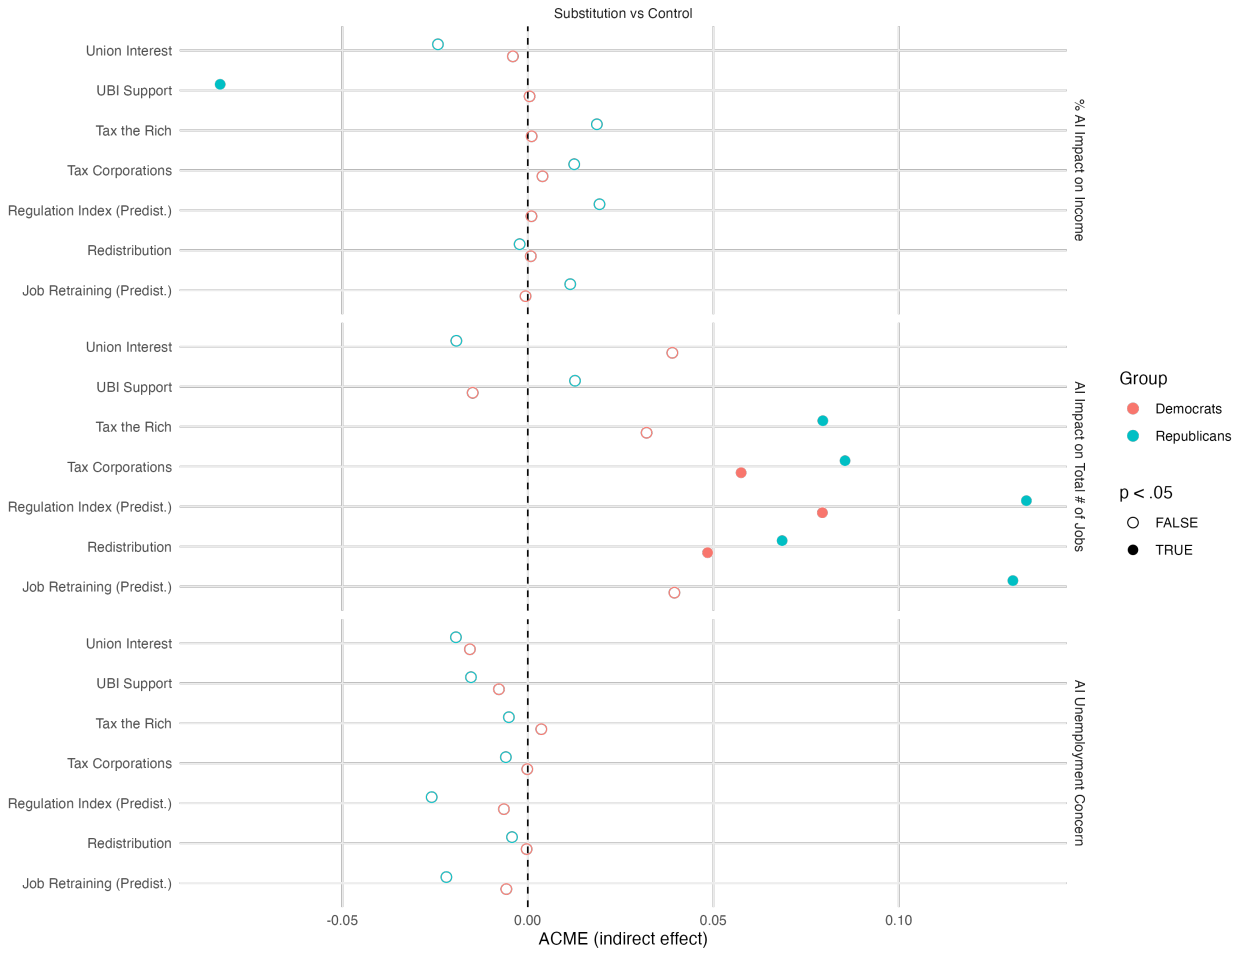


Figure 5: Democrats & Republicans, Substitution vs Control

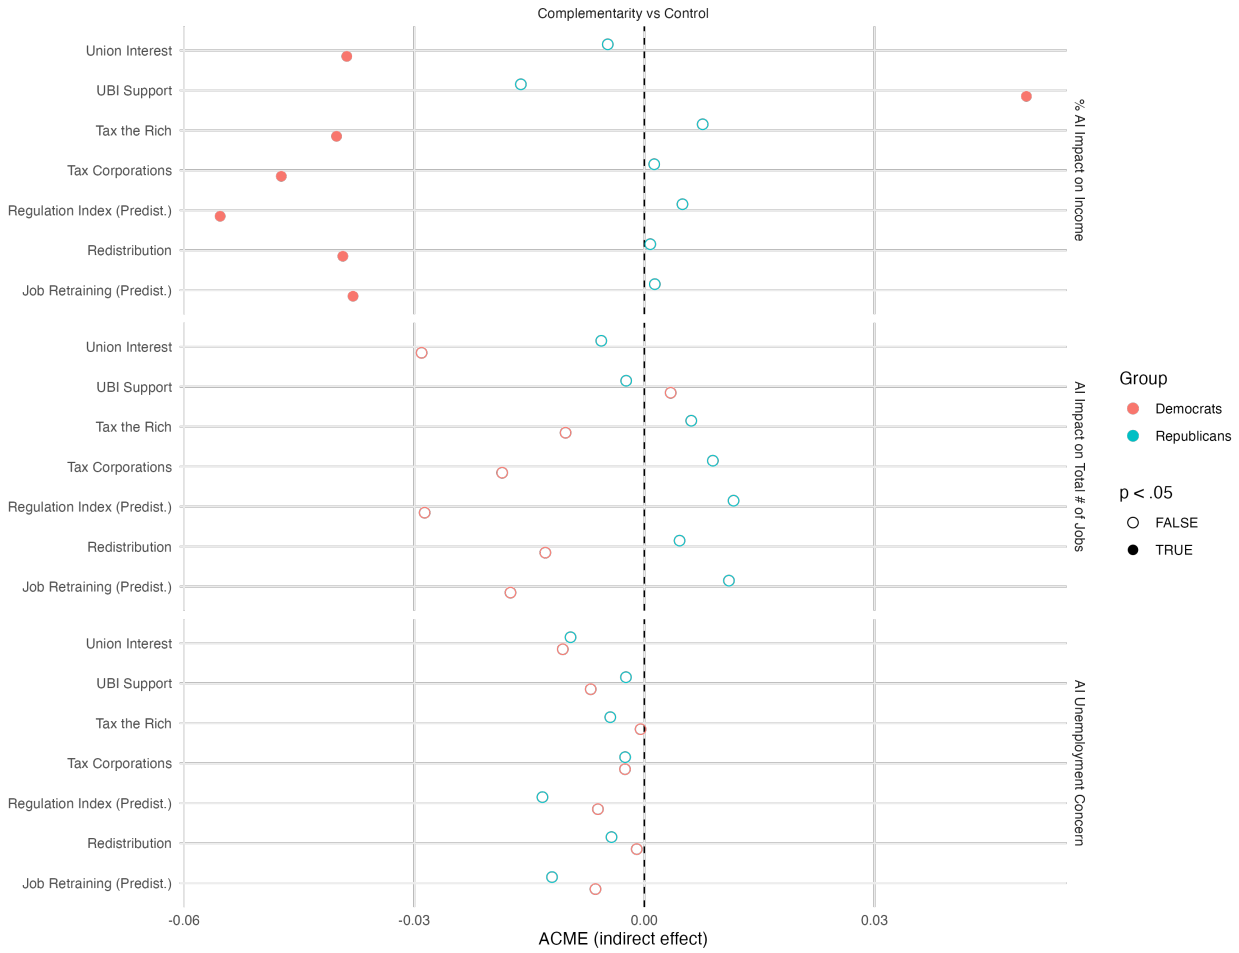


Figure 6: Democrats & Republicans, Complementarity vs Control

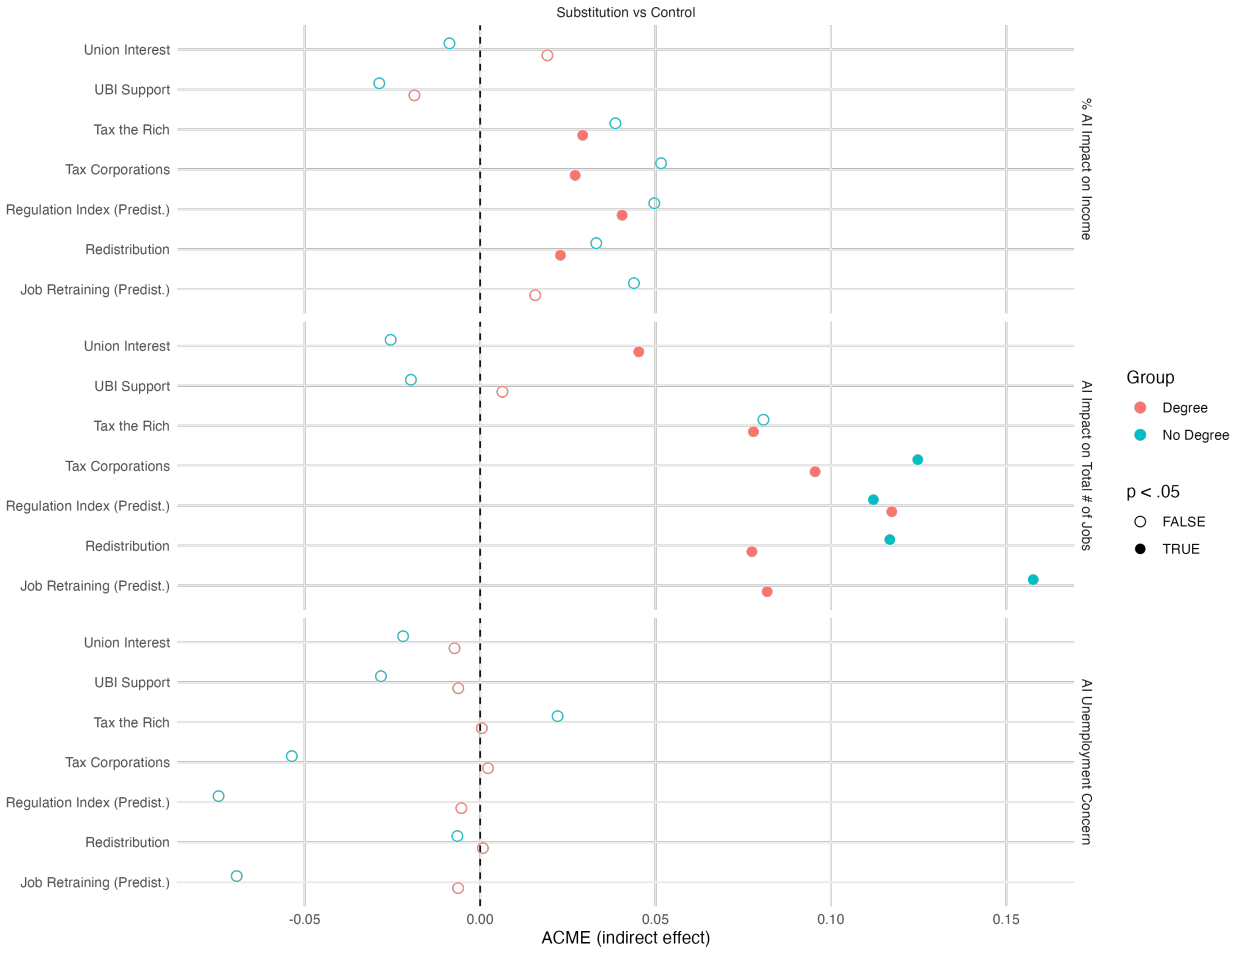


Figure 7: Degree & No Degree, Substitution vs Control

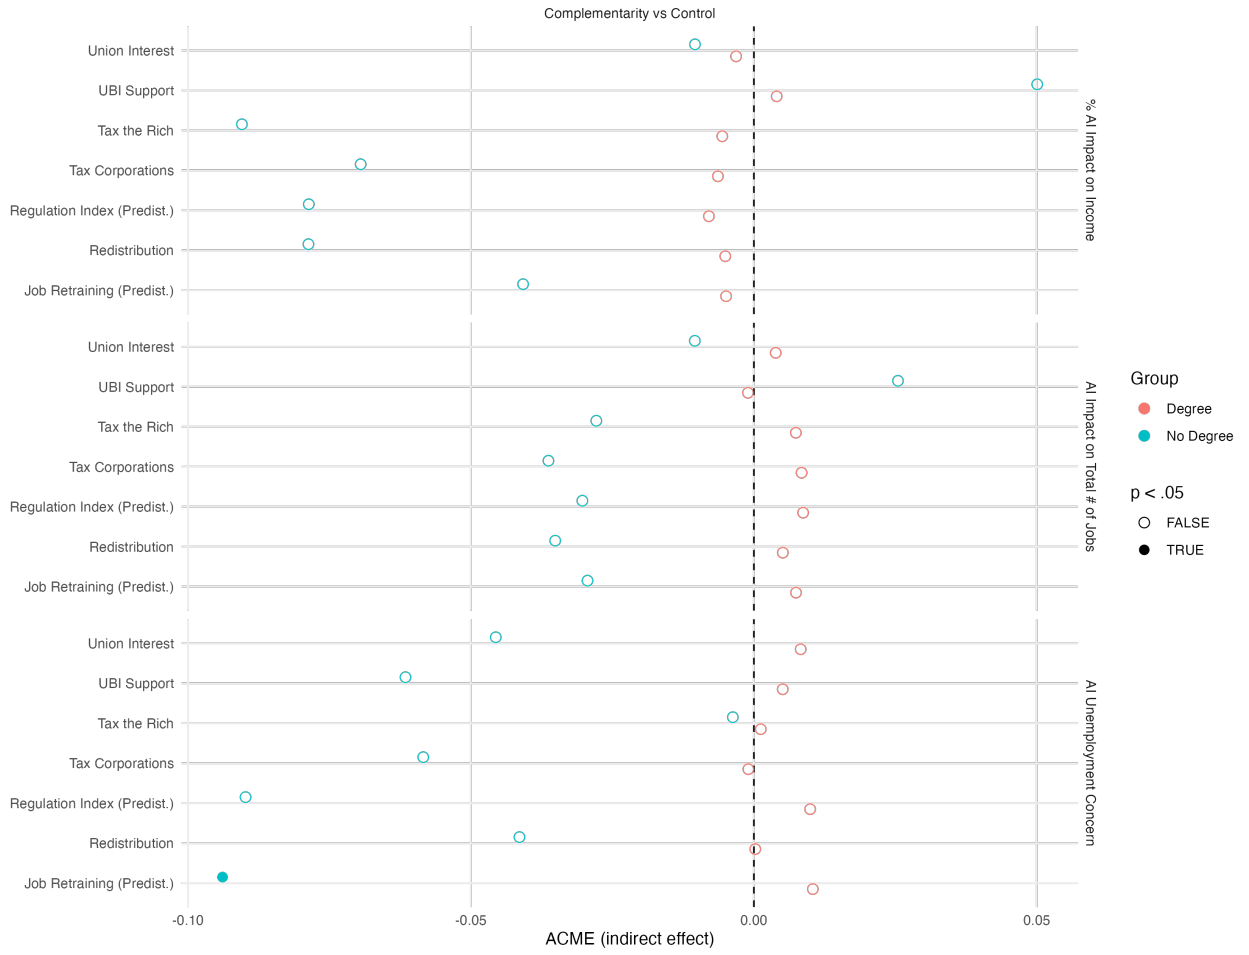


Figure 8: Degree & No Degree, Complementarity vs Control

G Other Analyses

G.1 Bonferroni Correction P Values

Following my pre-analysis plan, I applied Bonferroni corrections across all 15 outcomes (30 tests). After correction, only sociotropic AI concerns regarding AI's potential to reduce the total number of jobs in the Substitution condition returned a significant adjusted p-value in the main analysis (p.adj = 1.30e-06). Expected future income imapcts from AI returned an adjusted p-value of 0.41 and all remaining outcomes returned adjusted p-values of 1. I also conducted subgroup analyses. Of all subgroup × outcome combinations tested, only sociotropic AI concerns in the Substitution condition among Republicans (p.adj = 0.003) and degree-holders (p.adj = 5.73e-05) survived correction. All other subgroup results returned adjusted p-values above 0.05 or equal to 1. Full correction tables are available from upon request.

G.2 Placebo ANOVA Test

Here, I test whether any of the four placebo videos shown to the control group had effects on any outcomes of interest. As shown, none of the placebo videos have statistically significant effects on respondents' policy attitudes or economic forecasts.

Table G.1: Placebo Tests: Placebo Video Assignment and Policy Outcomes

Outcome	$\chi^2(3)$	p-value
Regulation Index (Predist.)	0.703	0.872
Redistribution	0.244	0.970
Union Interest	3.146	0.370
Job Retraining (Predist.)	0.941	0.816
Tax Corporations	1.234	0.745
Tax the Rich	1.717	0.633
Redistribution (normalized)	0.459	0.928
UBI Support	2.588	0.460
AI Unemployment Concern (normalized)	0.846	0.839
EFI AI (normalized)	2.882	0.410
Sociotropic AI Expectations (normalized)	5.391	0.145
Ban AI Layoffs	1.384	0.709
General AI Regulation	0.022	0.999
AI Safety Regulation	3.985	0.263
Create Federal AI Regulator	0.525	0.913

Note: † $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$. Robust Wald χ^2 tests (3 df) of joint significance of placebo video assignment on each z-scored outcome. Null results support the validity of randomization.