Automation and the Salience of Protectionism*

Zuhad Hai †

November 14, 2019

Working Paper

Abstract

Why does automation have fewer political opponents than international trade openness? In this paper, I argue that displacement due to automation gets politically attributed to trade. This mis-attribution occurs because those displaced by automation believe trade to be behind job loss. I test this claim by showing that displaced workers are more likely to erroneously file petitions for Trade Adjustment Assistance (TAA) in labor markets that have adopted industrial robots as a substitute for labor. Further, automation-exposed communities are more likely to correctly identify their political representative's position on trade (but not non-trade issues). I interpret this as evidence that automation-displaced communities believe a model of the economy where trade is to blame for job loss. Overall, the lack of compensation for automation-based displacement then might be the result of a political system that artificially raises the salience of trade policy and that therefore fails to accurately aggregate information about the causes of job loss.

*I would like to thank Judith Goldstein, Mike Tomz, Avidit Acharya, and Justin Grimmer for their encouragement and support in pursuing this project. This project was made better by feedback and help from them and from Feyaad Allie, Valentin Figueroa, Bobby Gulotty, Soyoung Lee, Rachel Lienesch, Sarah Thompson, and Cesar Vargas. All errors are my own.

†Stanford University. Email: zuhadhai@stanford.edu

1 Introduction

A puzzling feature of politics in the developed world in the past 30 years is the absence of an organized political movement advocating the restriction of automation: the adoption of production technologies that displace workers. This is puzzling, not just in the face of vast economic displacement due to automation, but also when contrasted with the robust movement against globalization that has arisen in developed countries in recent years.

Why does automation have fewer political opponents than international trade openness? I argue that since layoffs due to automation and trade might affect similar types of workers, workers displaced by either might look to each other to ascertain the cause of their job loss. In the presence of long standing institutions and policies around trade, automation affected communities are highly primed on trade policy and the remedies associated with its effects. In this case, automation exposed workers will believe trade to be behind their job loss. If these communities then blame trade, this will inflate the political reaction against international trade, and deflate that against automation.

I provide evidence for this claim by testing two of its observable implications. First, I show that local labor markets in the US that are more exposed to industrial robots have a higher likelihood of erroneously petitioning the government for Trade Adjustment Assistance (TAA). These same regions have both a higher rate of TAA petition denial and a higher absolute number of denied petitions per capita. This shows that automation displaced workers believe that their job loss is plausibly due to international trade[1].

Next, I show that automation affected communities also hold high levels of knowledge about *national* level trade policy, which is often thought of as less salient for voters than other economic issues. I use data from the 2006 Cooperative Congressional Election Study (CCES) to show that respondents in automation exposed commuting zones are more likely to correctly identify both their Senators' votes on the Central American Free Trade Agreement (CAFTA). Using similar placebo tests for non-trade issues, I find that this effect is unique to trade policy and is not indicative of voters being better informed of non-trade issues such as the minimum wage or abortion.

The results together imply that those exposed to automation believe a model of the economy where trade is to blame for job loss. Overall, I make progress in resolving the puzzle of why global-

[1]Throughout the paper, I mean international trade to encompass all aspects of international commercial openness that displace workers. This therefore includes forces such as off-shoring of production.

ization has more political opponents than automation: the political system is unable to accurately classify job loss as automation related because ill-informed workers have an incentive to represent their job loss as trade related.

This explanation stands in contrast to several alternatives. Most notably, my explanation is in contrast to the idea that laws against automation are considered less fair by voters (Rodrik, 2018), or that firms reliant on automation lobby against laws restricting automation. Although my explanation does not wholly exclude these alternatives, it provides an alternative that also helps explain the large backlash against globalization in highly developed economies.

The results presented here also have broad implications for understanding the politics of the world economy. In order to make the post-World war II project of globalization and trade-openness politically self-sustaining, governments have to compensate those displaced by trade. In this paper I argue that an important pre-requisite for such compensation is that that political system needs to be able to accurately separate trade-related displacement from non-trade-related displacement. At stake here is a self-sustaining project of international openness and a workforce that accurately represents its political grievances to politicians.

These findings also have broader implications for the political voice against globalization. The political reaction to globalization seems at odds with findings that those who are actually most hurt by globalization hold views not based on their self interest (Hainmueller and Hiscox, 2006) and do not punish politicians for going against voters' trade policy preference (Guisinger, 2009). Further, trade policy and institutions themselves might affect the causal beliefs of voters. Kim and Gulotty (2018) for example, show that trade adjustment assistance serves as a mediator that informs political actors such as voters of the size of displacement due to trade. Further, Congressional representatives get notified of TAA petitions filed in their constituency and their outcomes further sending signals to political actors that might misinform them about the size of trade-related displacement. Kim and Pelc (2017) find that applying for TAA is itself mediated by political representatives' framing of the program and of the nature of job loss. I argue that institutions designed around trade might affect the beliefs of displaced workers whose job loss was not due to trade.

Methodologically, this paper makes several contributions to IPE research. Firstly I rely on a behavioral measure of individual beliefs using TAA petitions. Since filing a TAA petition is costly and fraudulent petitions face legal sanctions, a denied petition provides evidence that the petitioner(s) believed trade to be behind their job loss. Further, unlike previous research, the use of TAA petitions allows me to observe actually displaced workers, rather than infer their level of economic insecurity or risk of job loss from their industry of employment and/or education level. Further, the survey based measure of voter knowledge about trade and non-trade Senate votes is linked to actual votes by the respondent's Senators, eliminating the risk of social desirability bias.

Before presenting the data and empirical strategy used in this paper, I first discuss the theoretical issues at stake, particularly why automation-displaced workers would attribute their job loss to trade. I then present my empirical analysis to test an implication of this theory: workers displaced by automation will be highly primed on remedies to trade-related displacement and national level trade policy. I interweave descriptions of the data used and the analyses themselves to ease the exposition of the argument. I then conclude with a synthesis of what the empirical analysis tells us about the politics of automation and trade, and what is left unexplained.

2 Automation and Trade in the United States

Why would automation-displaced workers attribute their job loss to trade? I argue that trade and automation have displaced similar types of workers, and therefore displaced workers will look to each other to ascertain the causes of their job loss. Since institutional structures and policy responses exist for trade, these workers will be highly primed on trade policy. This will lead them to publicly represent their job loss as trade-related, even if the true culprit lies in automation.

Theoretically, do we expect economically displaced workers to organize politically? My starting point is the dominant paradigm in international political economy, open-economy politics (OEP), in which those most hurt by an economic policy or economic phenomenon will form preferences based on their self-interest and expend resources in the political sphere to counteract this policy until the expected costs of doing so outweigh the benefits.[2]. This paradigm correctly predicts that the economic displacement due to import competition in the United States in the past 30 years would lead to a popular backlash against international trade[3]. However, applied to automation *within* a

[2]See Lake (2009) for an overview of the open-economy politics paradigm.

[3]Following a wave of recent work, the connection between trade-related displacement and anti-globalization collective action is by now well-established. Most of this work focuses on isolating the causal impact of China's rise in the international economy on politics in developed countries. Feigenbaum and Hall (2015), for example, find that local shocks to employment due to Chinese import competition leads Congressional representatives to become more protectionist in their roll-call voting, in order to avoid losing office. Autor et al. (2016a) find similarly that import exposure lead to the replacement of moderate legislators with more ideologically extreme ones. Further prominent work showing that trade exposure leads to political collective action, including in non-US contexts, is provided by Jensen et al. (2017), Dippel et al. (2015), Colantone and Stanig (2018) and Ballard-Rosa et al. (2017) among others.

country, OEP should also predict a complementary political voice demanding the regulation or restriction of automation[4].

Why have workers that were displaced by automation not formed a political movement opposing it? More realistically, why have political entrepreneurs not emerged who specifically target their policy proposals towards the regulation of automation or compensation of those most hurt by it?[5] One precondition for such political action is that those most hurt by automation be able to understand the cause of their job loss. If workers do not know that automation caused their job loss, they might blame trade, misrepresenting their interests to political leaders, who then might act as if their constituents demand compensation from trade-related job loss.

There is evidence to suggest that members of the public at large do not form causal beliefs about economic policy based on its impact on their own material well-being. Rho and Tomz (2017) argue that this is because people have very little information about how forces like trade affect their income. In the absence of 'economically correct' causal beliefs, there is evidence that people form beliefs *sociotropically*; based on the perceived impact of an economic policy or phenomenon on their community Mansfield and Mutz (2009). If automation-displaced workers form beliefs in this way, they will look to their community to ascertain the causes of economic displacement.

Recent work in political science has made progress on the question of why there is no backlash against automation. Specifically, while it is true that in countries like the United States, there is a political movement against globalization but not automation, recent work has shown that even those impacted by automation organize against globalization. Wu (2019) for example finds that

[4]Classical theories of collective action *do* outline several reasons why those most hurt by a policy or economic phenomenon might not be able to organize against it. Larger and more dispersed groups, for example, may find the expected costs of organizing to be greater than the expected benefits (Olson, 2009). Since automation in the United States has had a large labor displacing effect, it may be reasonable to attribute the lack of a political movement to group size. However, displacement due to automation in the United States has also been concentrated geographically. In the context of computerization, automation seems to be concentrated in urban areas (Autor et al., 2015), while in the context of industrial robots, there are clusters of heavy automation in places that specialize in automotive sector, such as Michigan (Acemoglu and Restrepo, 2017). I therefore focus on informational constraints to collective action since they seem to be most relevant here.

[5]In recent years, especially within the context of the 2020 US presidential election, automation-related job loss has indeed increased in salience. However, much of this concern seems forward-looking, targeting *anticipated* job loss due to to artificial intelligence, rather than compensating existing and past job loss due to the automation of lower skilled jobs. Notably, unlike the automation of the 1990s and 2000s analyzed in this paper, artificial intelligence displaces higher skilled workers, and therefore the theory outlined here does not naturally apply to it.

survey respondents at higher risk of automation tend to express political views against off-shoring and immigrants. Kaihovaara and Im (2018) similarly find that those employed in routine jobs (i.e. those at highest risk of automation) tend to hold highly negative views about immigrants. Rodrik (2018) similarly attributes this blame attribution to trade on the perceived 'fairness' of automation and 'unfairness' of trade, immigration and off-shoring. Anelli et al. (2019) and Im et al. (2019) find that exposure to automation is associated with a wholesale shift towards right-wing and ethno-nationalist parties. Further, these policy preferences are highly persistent. Zhang (2019) find that providing respondents with accurate information about the labor market effects of automation does not change their policy positions on right-wing policies on welfare, immigration and trade. Gingrich (2019) similarly finds that welfare spending by governments did not significantly moderate the effect of automation on vote for populist leaders.

As Figure 1 shows, these findings can be represented as the finding that the 'shock' of automation leads people to hold extreme policy preferences such as protectionism, nativism and right-wing ethno-nationalism. What is unclear, however, is whether these extreme policy preferences are simply a result of a dissatisfaction with the political status quo. Sociotropic, pocket-book and/or retrospective concerns could all lead automation displaced communities to hold extreme policy positions as a reaction to the status quo. This is important to know because if extreme policy position are simply the results of a dissatisfaction with the status quo, then it is still not clear why those exposed to automation don't demand policies restricting automation specifically.

Figure 1: **Theoretical Contribution:** Exposure to automation could lead to extreme policy preferences through many channels. In this paper I test the mechanism that automation exposure leads to protectionist preferences because automation exposed populations believe trade to be the cause of their job loss.

My contribution to this literature is to identify a novel channel by which job losses get misrepresented: those exposed to automation *believe* trade to be behind job loss. This explanation is complementary to, but distinct from, those based on cognitive biases, framing effects and/or fairness concerns. If trade and automation displaced workers are drawn from the same skill-distribution as trade-displaced workers, they might use actions and remedies sought out by trade-displaced workers as a reference point for their own collective action.

Do these two sets of workers actually look the same in terms of their skill-level? Theoretically, two separate bodies of work in economics tell us that automation and import competition in developed countries can end up displacing very similar types of workers. Specifically, work in labor economics has found that the biggest losers from automation tend to be middle and lower skilled manufacturing workers, since recent production technologies have been 'skill-biased' in the sense of being complementary to high skilled workers and displacing lower skilled workers. In the theoretical literature on international trade, the famed Heckscher-Ohlin theorem predicts that trade openness in relatively low-skill scarce countries such as the US will hurt low-skilled workers the most. Empirically, Autor et al. (2016b) find that trade-related job losses do concentrate highly in lower-skill manufacturing regions, while Acemoglu and Restrepo (2017) find that the effects on industrial robots on employment is concentrated in lower-skilled occupations.

Of-course, automation and trade related displacement need not be mutually exclusive. For example, a firm might automate in response to increasing price pressures due to foreign import competition. In these cases, we would expect there to be blame directed towards trade. However, it is difficult to make the case that the majority of job losses due to automation had import competition as their cause. In the United States, for example, automation of routine-intensive tasks is attributed mainly to the large influx of skilled workers who could complement the introduction of computers in the workplace (Acemoglu, 2002). Similarly, although the adoption of industrial robots does seem to be a function of international industrial trends, it is reasonable to expect that to have happened even in the absence of import competition. In fact Hicks and Devaraj (2015) estimate that only 13.4% of job losses in the United States during the period 2000-2010 were due to import competition, with the remainder attributable to productivity growth. Nevertheless, I do not *a priori* rule out that automation related displacement might be trade related. In my empirical analysis however, I will be able to isolate variation in the adoption of robots that is purely due to the pressures of the global technological frontier, mitigating this concern. On the dependent variable side, I will further be able to isolate job losses due to non-import competition (or off-shoring) by using petitions of Trade Adjustment Assistance (TAA) that are denied by the government.

Isolating the impact of automation is difficult because of the variety of types of displacement that may reasonably be labeled as automation. In this paper, I define automation as the replacement of labor with machinery for a given set of tasks that are an input to production. In many cases, such replacement may create demand for labor that might be employed in different tasks that are complementary to the new technology[6]. Although this definition requires the adoption of machinery in place of labor, it does not necessarily rule out the adoption of more efficient management practices (such as 'just-in-time' manufacturing). This is because such practices might also result in the adoption of advanced technology such as scheduling software, having a very similar employment effect as the adoption of spreadsheet software. For the purpose of measurement, in this paper I focus on a relatively stark form of automation: the adoption of industrial robots in place of manufacturing workers. Although this means that my empirical analyses do not directly capture the effect of related forms of automation such as computer software, it does provide me with a clear-cut case of labor being unambiguously replaced with machinery. To the extent that

[6]For example, installing an industrial robot might increase the demand for a technician who understands how to program and repair the robot.

computerization and other forms of automation displace similar types of workers, I expect the results to apply to those settings as well.

With an idea of how this project fits into the larger conversation about automation and globalization, I can state the central claim of this paper:

Claim: I argue that communities exposed to automation-related displacement believe that trade is behind their displacement. I derive two observable implications of the theoretical discussion above:

Observable Implication 1: *Following a job loss, automation exposed workers will erroneously apply for (and often get rejected for) assistance that is meant for trade-related job loss.*

Just because automation displaces workers rely on remedies for trade does not necessarily imply that this will translate into them having stronger opinions about international trade policy itself. Therefore, another test of the theory is whether automation exposed workers display high levels of knowledge about national trade policy:

Observable Implication 2: *Automation exposed workers will know more about national level trade policy than non-automation exposed workers, but this will not be true for non-trade policy.*

We would not expect these to be true if exposure to automation simply leads people to dislike the status quo and/or if those most hurt are averse to regulating technological change. I now move to operationalization these observable implications and testing them empirically.

3 Measurement and Empirical Analysis

3.1 Measuring Automation Exposure Using Robot Purchases by Industry

In order to measure exposure to automation at the local labor market level[7], I use the data compiled by Acemoglu and Restrepo (2017), who construct a measure of local labor market exposure to

[7] To define a local labor market, I use US commuting zones, which are regions of the US determined by the US Bureau of the Census to share a labor market, calculated using commuting times. I use the 741 commuting zones defined in 1990, but have consistent data on automation exposure on only 722 of them (Alaska and Hawaii are excluded, leaving 48 states).

industrial robots using data on purchases of industrial robots by countries around the world. They rely on a database compiled by the International Federation for Robotics (IFR), which collects yearly survey data from robot manufacturers (IFR, 2017).[8]

The starting point for constructing a regional automation exposure variable is to derive a measure of industry level exposure to automation. The industry level variable I use is roughly a measure of the change in the stock of industrial robots per worker for an industry over a given time period. More formally, for the US, over the time period 2004-2007, industry i exposure to robots is defined as:

$$Exposure\ to\ Robots : US_{\ i,2004-2007} = \left(\frac{M^{US}_{i,2007} - M^{US}_{i,2004}}{L^{US}_{i,1990}} - g^{US}_{i,2004-2007} \frac{M^{US}_{i,2004}}{L^{US}_{i,1990}} \right)$$

Where $M^{US}_{i,t}$ denotes the stock of industrial robots in industry i in year t, and $L^{US}_{i,t}$ is the number of workers employed in industry i in year t. $g^{US}_{i,2004-2007}$ denotes the growth rate of output in industry i during the time period $2004-2007$[9]. Data on growth rates as well as industry employment is from the EU KLEMS database (Jäger, 2017).

Following Acemoglu and Restrepo (2017), I use the adoption of industrial robots in 5 European countries (Denmark, Finland, France, Italy, and Sweden) to construct the instrument for adoption of industrial robots in the US. As the authors point out, these 5 countries were selected because of their robot adoption being far ahead of the United States. Germany was excluded despite also being ahead of the US because its robot adoption trends are so far ahead of other countries that it might not be relevant for explaining US adoption trends. The analogous industry level measure for these European countries is then:

$$Exposure\ to\ Robots : Europe_{\ i,1993-2007} = \frac{1}{5} \sum_{j \in Europe5} \left(\frac{M^{j}_{i,2007} - M^{j}_{i,1993}}{L^{j}_{i,1990}} - g^{j}_{i,1993-2007} \frac{M^{j}_{i,1993}}{L^{j}_{i,1990}} \right)$$

Where j indexes country, and the measure is then the average of the industry level exposure across the 5 selected countries. Table 1 presents the value of these variables across the different

[8]Following the International Organization for Standardization (ISO), the IFR defines an industrial robot as an "automatically controlled, re-programmable multipurpose manipulator programmable in three or more axes". This definition excludes other technological changes such as widespread computerization and the adoption of labor saving management practices such as 'just-in-time' manufacturing. Although these could also fall under the definition of automation, I only look at the adoption of industrial robots as it has a clear labor-displacing effect Acemoglu and Restrepo (2019).

[9]Ideally we would use the two time periods 1990-2000 and 2000-2007. However, the IFR data do not show disaggregated industry shares for the US prior to 2004, so we use the time period 2004-2007.

industries reported by the IFR. While utilities and services predictably rank lowest in terms of industrial robot adoption, the automotive industry features by far the highest degree of industrial robot adoption in the time period, both in Europe and the US. This accords with the auto industry featuring prominently in media and other public discussions of automation[10].

	IFR Industry Name	$US\ Exposure\ to\ Robots_{i,2004-2007}$	$Euro5\ Exposure\ to\ Robots_{i,1993-2007}$
1	Automotive	17.817	32.936
2	Electronics	2.200	3.457
3	Plastics and Chemicals	1.593	21.464
4	Basic Metals	1.106	5.701
5	Food and Beverages	0.899	5.202
6	Metal Products	0.664	8.009
7	Miscellaneous	0.586	-1.202
8	Industrial Machinery	0.195	1.005
9	Minerals	0.111	2.660
10	Shipbuilding and Aerospace	0.060	2.827
11	Wood and Furniture	0.007	3.649
12	Textiles	0.005	1.060
13	Construction	0.004	0.070
14	Mining	0.004	2.687
15	Agriculture	0.003	0.159
16	Education and Research	0.003	0.295
17	Paper and Printing	0.002	0.614
18	Services	0.000	0.000
19	Utilities	0.000	0.022

Table 1: Industries by Robot Adoption in the US and the 5 European countries, arranged in descending order of US adoption. The Industry classifications follow the IFR industry names.

3.1.1 Shift-Share Variable Construction

The main commuting zone level variable for automation exposure is constructed as follows. As with all shift-share variables, the exposure variable into two parts: an industry level 'shift' variable, and an industry-region level 'share' variable that distributes the industry level shift across regions (Adão et al., 2019). A typical shift-share variable then has the form:

$$Exposure_{c,t_0-t_1} = \sum_i Share_{i,c,t0} \mathbf{X} Shift_{i,t_0-t_1}$$

[10] https://advancedmanufacturing.org/automotive-industry-improves-automation/ https://abc7chicago.com/automotive/ford-unveils-$1-billion-upgrade-at-chicago-plants-adds-jobs/5362194/

Where c indexes region, i indexes countries, and t_0 and t_1 are the beginning and end years of the time period under consideration. In many setting (such as ours), this variable is then instrumented using an analogous variable of the form:

$$Exposure\ Instrument_{c,t_0-t_1} = \sum_i Share_{i,c,t_{-1}} \mathbf{X} Shift\ Instrument_{i,t_0-t_1}$$

Where the share variable $Share_{i,c,t_{-1}}$ is sufficiently lagged, and a shift variable is selected $Shift\ Instrument$ that satisfies the exclusion restriction. That is, we expect the causal effect of $Shift\ Instrument$ on our dependent variable to operate only through its effect on $Shift$. In this setting, as I explain below, the main shift-share variable I will use is a measure of adoption of robots by US commuting zone, and then instrument this using lagged employment shares and a shifter that measures the adoption of industrial robots by countries that are similar to the US but further ahead on the global robotics adoption frontier.

3.1.2 Main Explanatory Variable: US Exposure to Robots

In this case, the exposure of a local labor market (specifically a commuting zone) to automation from 2004 to 2007 is calculated as an average of nation-wide industry level exposures, weighted by a variable denoting the share of the commuting zone's economic activity in each industry:

$$Exposure\ to\ Robots : US_{c,2004-2007} = \sum_{i \in Industries} l_{c,i,1990} \mathbf{X} Exposure\ to\ Robots : US_{i,2004-2007}$$

The weighting variable used, $l_{c,i,t}$, is a commuting zone's 1990 share of employment in industry i. Intuitively, we calculate the commuting zone's exposure to robots by averaging exposure across all industries in that commuting zone, weighting by how much of the commuting zone's labor force is employed in an industry.

3.1.3 Instrument: European Exposure to Robots

As before, I follow the instrumental variable strategy used by Acemoglu and Restrepo (2017) and use the industry level exposures to industrial robots in the five selected European countries.

More specifically, the industry level exposure instrument is:

$$Exposure\ to\ Robots : Europe_{c,1993-2007} = \sum_{i \in Industries} l_{c,i,1970} \mathbf{X} Exposure\ to\ Robots : Europe_{i,1993-2007}$$

Where l_{ci}^{1970} is the 1970 share of commuting zone c's labor force in industry i. This weighting variable is chosen to be from a time period sufficiently in the past (before the onset of widespread adoption of industrial robots in the US) to avoid confounding the effect of the 1990 share.

Figure 2 shows the geographic distribution of the predicted exposure to automation ($Exposure\ to\ Robots : US_{c,2004-2007}$ as predicted by $Exposure\ to\ Robots : Europe_{c,1993-2007}$) and compares it to the predicted change in imports per worker from China in the same time period (the so called 'China Shock') from Autor et al. (2015)[11]. The automation shock from industrial robots seems to be heavily concentrated in the US Midwest, especially Michigan. This accords with the industry level exposure being highest for the automotive industry, which has historically been concentrated in Michigan, especially the Detroit metropolitan area.

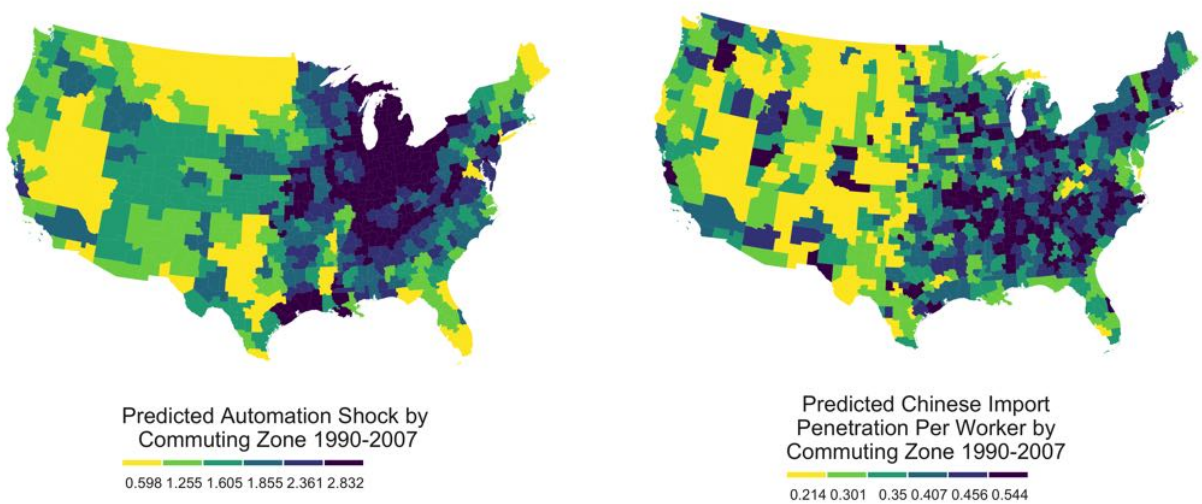

Figure 2: Automation vs Trade Exposure by Commuting Zone: Predicted Change in Industrial Robot Adoption from Acemoglu and Restrepo (2017) and Predicted Change in Chinese Import Penetration from Autor et al. (2015)

[11]This variable, called the change in import penetration per worker or $\Delta IPW_c^{China-US}$ is also a shift-share style variable, which we use predicted values of by regressing it on $\Delta IPW_c^{China-Other}$, a similar shift-share instrument. Since it is not the focus of the paper, I will not go into detail about its construction. It is constructed similarly to the automation exposure variable by averaging over change in Chinese imports per worker in an industry over all industries, weighting by commuting zone shares of that industry: $\Delta IPW_{c,1990-2007}^{China-US} = \sum_j l_{c,i,1990}$ **x** $\frac{\Delta Imports_{i,1990-2007}^{China-US}}{L_{j,1990-2007}}$. This is instrumented using an analogous variable where the share is lagged 10 years and the change in Chinese imports is for 8 other advanced economies similar to the US. I then use the predicted values of $\Delta IPW_{c,1990-2007}^{China-US}$ in this paper.

3.1.4 First Stage and the Exclusion Restriction

The identification strategy relies on some countries being 'leaders' and others being 'followers' in adopting technologies from the global technological frontier. Our instrument will be valid to the extent that the five European countries' adoption of robots is a result largely of domestic factors within those countries and affects the salience and awareness of trade policy among US workers only through its effects on US robots adoption.

There is ample evidence that the US is behind several European countries in the innovative frontier in industrial robot development and adoption. As the Wall Street Journal reported in 2017: "A report to President Barack Obama on advanced manufacturing, prepared by his council of science advisers in 2012, concluded that the "hard truth" was that the U.S. lagged other rich nations on manufacturing innovation." (Michaels, 2017). Further, the IFR data itself shows the US consistently lagging behind the five selected European countries. The graph below, replicated from Figure 1 in Acemoglu and Restrepo (2017), shows that the stock of industrial robots per 1000 workers was consistently lower in the US than in the selected 5 European countries.

Figure 3: The US Lags behind the 'Euro 5' in Robot Adoption

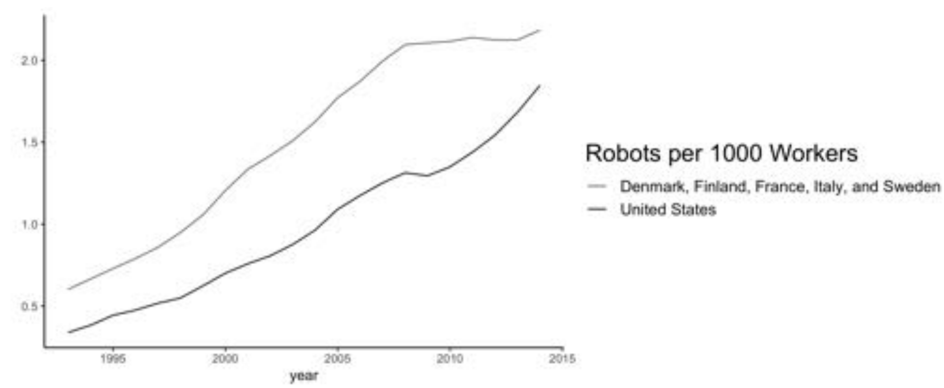

Why did countries in Europe adopt robots faster than the US? One leading reason is that they had a much faster aging population than the US. The link between aging and automation is well documented and often mentioned in media reports of automation (Acemoglu and Restrepo, 2018; Eglitis and Seputyte, 2018). Firms in countries with aging populations have a greater incentive to automate jobs that middle aged workers have a comparative advantage in, which are typically related to production (Acemoglu and Restrepo, 2018). Since the US is atypical among developed countries for having a younger population, we expect automation trends in the US to be below those of rapidly aging countries in Europe.

Further, it seems plausible that countries with initial leadership in automation technology continued to build on this advantage through competitive industry pressure and government funding. The same Wall Street Journal article also states: "In Japan and Europe, industries such as electronics and pharmaceuticals pushed their automation suppliers for increasingly specialized equipment. Governments funded research and development." This gives confidence to the claim that European adoption of robots causally affects automation-displaced workers in the US only through its effect on US adoption of robots.

Table 2 documents a strong first stage effect, at the level of US commuting zone, of a regression of $Exposure\ to\ Robots : US_{c,2004-2007}$ on $Exposure\ to\ Robots : Europe_{c,1993-2007}$. As column 3 shows, a 1 unit change in $Exposure\ to\ Robots : Europe_{c,1993-2007}$ is associated with a 0.212 unit change in $Exposure\ to\ Robots : US_{c,2004-2007}$, corresponding to roughly 3/4th standard deviation increase in $Exposure\ to\ Robots : US_{c,2004-2007}$.

Table 2: **First Stage**

	(1)	(2)	(3)
	Dependent Variable: *Exposure to Robots: US*		
Exposure to robots: Europe	0.214***	0.188***	0.188***
	(0.024)	(0.026)	(0.026)
Dep Var Mean	0.330	0.330	0.330
R squared	0.553	0.628	0.631
Total Petitions Control?	No	No	Yes
Trade Shock Control?	No	Yes	Yes
# States	48	48	48
# Czones	722	722	722

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Standard errors are clustered at the state level and reported in parentheses. The unit of analysis is the commuting zone (CZ) - decade. For all per-capita dependent variables, the denominator is commuting zone population in 1990.

Now that I have a measure of automation exposure at the level of the local labor market, I can relate that to measures of workers beliefs about the causes of their job loss. In the next section, I describe a measure of workers mis-categorization of their job loss and empirically relate it to the exposure of labor markets to automation.

3.2 Observable Implication 1: Automation Displaced Workers Will File for Trade Adjustment Assistance (TAA)

Measuring Mis-categorization of Job Loss: I use the filing of petitions for Trade Adjustment Assistance (TAA) by displaced workers to develop a measure of displaced workers belief that their job loss is trade-related. Trade Adjustment Assistance (TAA) is a Federal program developed in the context of widespread trade liberalization in the second half of the 20th century. Instituted under the Trade Act of 1974, the largest part of the TAA program is designed for displaced workers and administered by the Department of Labor's Employment and Training Administration (US Department of Labor Employment and Training Administration, 2018).

Under TAA, upto a year after a layoff, workers may file a petition to be compensated for their job loss, as well as obtain retraining benefits. Importantly, workers have to present the case for their job loss being related to international trade. This could be if their job loss was caused by an increase in import competition or if production in their firm has been (or is being) shifted to a foreign country. Workers may also file a petition if they believe that their job loss is due to a loss of business to a TAA-certified firm. Until 2002, the TAA program covered only 'production' workers, which meant that only workers laid off in the manufacturing sector could file a petition. In 2002, service workers were made eligible to apply to TAA, and job losses due to the outsourcing of services were added. Since industrial robot adoption is not relevant for the service industry, I exclude petitions from the service sector in all analyses.

Measure Validity Under what assumptions does filing an eventually denied petition reflect a belief that trade is behind job loss? Suppose that a group of workers that have been laid off are deciding whether or not to file a TAA petition. The workers know that they have to make a plausible case for why they think that their job loss is trade-related. Let p denote their belief that trade is behind their job loss. Suppose workers believe that the TAA administration correctly ascertains the cost of their job loss, and so p is also the probability that the petition will be approved, if filed. If worker's file, their expected payoff is

$$EU[filing] = p * Benefits - Cost$$

while their expected payoff from not filing a petition is 0. The group of workers will then file a petition if the expected benefits from doing so exceed the expected net benefits of not filing:

File a Petition If:

$$EU[filing] \geq EU[Notfiling]$$
$$\implies p * Benefits - Cost \geq 0$$
$$\implies p \geq \frac{Cost}{Benefits}$$

Suppose a researcher compares two labor markets with equal numbers of displaced workers, and equal amount of exposure to trade, but with one labor market having a larger share of its job losses due to automation. If the more automation exposed labor market features a higher rate of TAA petition denial, this could be the result of three relationships (not necessarily mutually exclusive). Firstly, it could be that automation exposure increases p: workers' perceived belief that trade is behind their job loss. Secondly, it could be that automation exposure increases the value of the benefits from getting certified for trade adjustment assistance. Thirdly, it could be that automation exposure decreases the cost of filing a TAA petition.

In the following analysis, I will assume that only the first of these three are true. That is, I assume that an increase in denied TAA filings following exposure to automation is indicative of a change in belief about the cause of job loss, rather than a change in the costs or benefits of filing a TAA application.

> *Assumption #1: Automation displaced workers do not value the benefits from certification more than non-automation displaced workers.*

In terms of benefits from getting certified, even though TAA benefits are nominally the same for trade versus automation exposed workers, these two sets of workers might value these benefits differently. I assume that this is not true for two main reasons. First, in the analysis to follow, I control for labor market characteristics like the change in the employment to population ratio, which might make the value of unemployment benefits different in automation vs non-automation exposed labor markets. Second, I rely on the idea that automation and trade displaced production workers come from the same quantile of the skill distribution and therefore face a similar job market.

> *Assumption #2: Automation displaced workers do not perceive the cost of filing a TAA petition to be lower than non-automation displaced workers.*

The assumption that automation exposed workers don't face a lower cost to filing a TAA petition is realistic if we realize that automation exposed workers have to work harder to make a case for why their job loss is trade-related. This might mean that the cost of filing a petition is actually *higher* for automation as opposed to trade-displaced workers.

With these two assumptions, an increase in denied TAA filings following automation exposure will be a result of a change in p: the workers belief that trade is behind job loss. I collect petition-level data from US Department of Labor Employment and Training Administration (2018) and subset to the period 2004-2007 inclusive. I then aggregate the petitions to the commuting zone level using county identifiers in the petition level data[12]. I create four variables related to TAA filing activity at the commuting zone level: the number of TAA petitions (services excluded) per 10,000 population filed in a commuting zone in the given time period, analogously the number of denied and certified petitions per 10,000 people filed in a commuting zone and the percent of petitions denied in a commuting zone in the given time period. Table 3 presents some summary statistics for these variables, alongside summary statistics for the automation related variables mentioned earlier.

Table 3: **Summary Statistics for Commuting Zone Level Variables**

	N	Mean	Std.Dev	Min	Max
Denied Petitions per 10,000 People	739	0.91	1.74	0.00	13.01
Certified Petitions per 10,000 People	739	2.62	4.12	0.00	44.88
Total Petitions in 2000s	739	11.40	27.68	0.00	313.00
% of Petitions Denied	739	0.16	0.23	0.00	1.00
Exposure to Robots: US	722	0.33	0.31	0.04	2.65
Exposure to Robots: Europe	722	1.63	1.04	0.35	9.04
Predicted Change in Chinese Imports per Worker	722	0.46	0.14	0.05	1.90
Change in Mexican Imports per Worker	722	0.99	1.28	-0.01	19.26
Change in Employment to Population Ratio, 1990-2008	722	2.63	2.65	-6.28	14.50

Notes: The Trade Adjustment Assistance (TAA) variables are all from the time period 2001-2007 inclusive.

The average denial rate by commuting zone was 16% of petitions filed in this time period. However, 2,017 out of the total 8,689 (23%) petition filed in this time period were denied, the higher

[12] I use a county to 1990 commuting zone crosswalk developed by The Health Inequality Project (https://healthinequality.org/data/)

rate being indicative of some commuting zones having a higher denial rate than others. The map in figure 4 shows the geographic distribution of TAA petitions across commuting zones. The darkness of each dot represents the denial rate of petitions in a commuting zone, and the size denoted the raw number of petitions by commuting zone (not per capita).

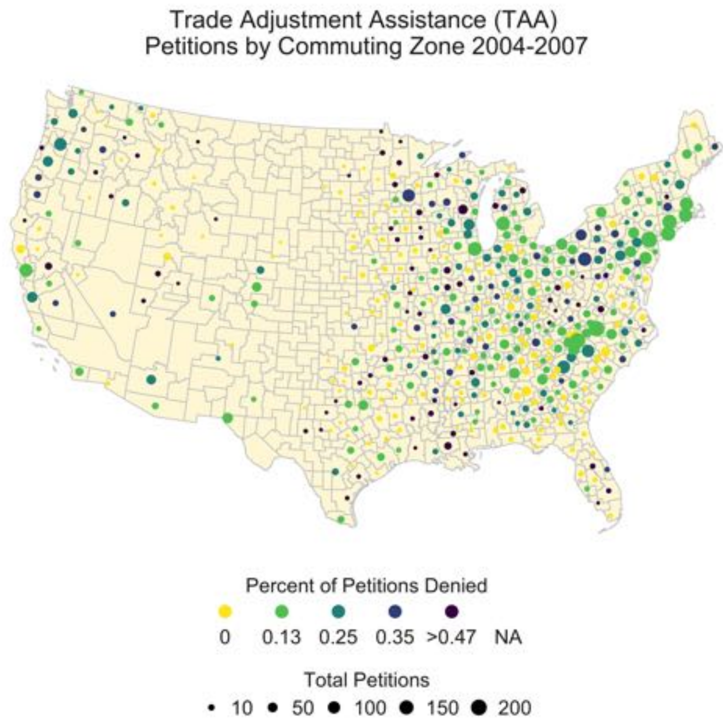

Figure 4

A visual comparison with figure 2 suggests that commuting zones with higher predicted exposure to industrial robots have a higher rate of TAA denial, denoted by darker dots. I now turn to testing this suggested relationship more formally.

3.2.1 Design

In order to test whether patterns of TAA petitions are systematically different in automation exposed commuting zones, I estimate two-stage least squares instrumental variables regressions of the following form:

$$\text{First Stage}:$$
$$Exposure\ to\ Robots: US_{c,2004-2007} = \beta_0 + \beta_1 Exposure\ to\ Robots: Europe_{c,1993-2007} \quad + \beta \mathbf{X}_c + \epsilon$$
$$\text{Second Stage}:$$
$$TAA_c = \alpha_0 + \alpha_1 Exposure\ to\ \widehat{Robots}: US_{c,2004-2007} \qquad + \beta \mathbf{X}_c + \epsilon$$

Where the second stage dependent variable TAA_c is one of the four TAA-related variables defined above. $\mathbf{X}_c$ is a matrix of control variables. Depending on the specification, the three control variables I use are: (1) the predicted change in imports from China per worker in a commuting zone from 1990-2007 from Autor et al. (2015), (2) the actual change in Mexican import per worker by commuting zone from Acemoglu and Restrepo (2017), and (3) the natural log of the total number of TAA petitions filed per 10,000 people in a commuting zone. Additionally, all regressions control for the change in the employment to population ratio of a commuting zone from 1990 to 2007, from Acemoglu and Restrepo (2017). The regressions are cross-sectional, with identification coming from variation in petition activity and automation exposure *across* commuting zones.

3.2.2 Results

The results from the first stage were presented in table 2. The results from the second stage are presented in table 4. Column 1 presents the result without control variables, while columns 2 and 3 present results when progressively adding trade shocks and the total number of petitions as controls, respectively. The total number of petitions control only makes sense with the dependent variable being the percent of petitions denied, so it is only included for those regressions.

Table 4: **Automation-Exposed Commuting Zones Have More TAA Petitions**

	(1)	(2)	(3)
Panel A: Effect on # Petitions Filed	Dependent Variable: *TAA Petitions per 10,000 People*		
Exposure to robots: US	3.078**	1.325	
	(1.328)	(1.687)	
Dep Var Mean	3.900	3.900	
# Czones	722	722	
R squared	0.160	0.238	
Panel B: Effect on # Denied Petitions	Dependent Variable: *Denied Petitions per 10,000 People*		
Exposure to robots: US	1.114***	0.839*	
	(0.387)	(0.490)	
Dep Var Mean	0.907	0.907	
# Czones	722	722	
R squared	0.051	0.080	
Panel C: Effect on # Certified Petitions	Dependent Variable: *Certified Petitions per 10,000 People*		
Exposure to robots: US	1.619*	0.268	
	(0.902)	(1.088)	
Dep Var Mean	2.621	2.621	
# Czones	722	722	
R squared	0.159	0.244	
Panel D: Effect on Denial Rate	Dependent Variable: *% of Petitions Denied*		
Exposure to robots: US	0.135***	0.131***	0.113***
	(0.042)	(0.045)	(0.044)
Dep Var Mean	0.161	0.161	0.161
R squared	0.010	0.014	0.025
Total Petitions Control?	No	No	Yes
Trade Shock Control?	No	Yes	Yes
# States	48	48	48
# Czones	722	722	722

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Standard errors are clustered at the state level and reported in parentheses. The unit of analysis is the commuting zone (CZ) - decade. Panel D also controls for the natural log of total petitions in the commuting zone-decade. All regressions exclude TAA petitions in the service sector. All regressions control for the change in the employment to population ratio of a commuting zone from 1990 to 2008. For all per-capita dependent variables, the denominator is commuting zone population in 1990.

Panel A presents the estimated effect of automation exposure on the total number of petitions filed in a commuting zone. Column 2 shows that after controlling for import competition from China and Mexico, the effect of automation on filing TAA petitions is large and positive, but statistically indistinguishable from zero. Ofcourse the total number of petitions combines denied and certified petitions, so the next two panels present the results for these two sets of petitions separately.

Panel B presents the estimated effect of automation exposure on the number of *denied* petitions filed in a commuting zone. Column 2 shows that after controlling for import competition from China and Mexico, the effect of automation on filing eventually denied TAA petitions is large, positive and statistically distinguishable from zero. To put the coefficient in perspective, an increase in automation exposure from the 10th percentile to the 90th percentile (an increase of 0.53 units of automation exposure) caused an estimated increase of 0.44 petitions per 10,000 people, roughly a 49% increase over the average of 0.907 denied petitions filed per 10,000 people.

I similarly interpret the estimated effect on the number of certified petitions. This effect, shown in Panel C, is positive but quote noisy, and fails to reach conventional levels of statistical significance after controlling for trade exposure from China and Mexico.

In panel D I test whether the *rate* of petition denial is affected by an exposure to automation. Column 3 shows that after controlling for import exposure and the total number of petitions filed in the commuting zone, automation exposure has a positive and statistically significant effect on the rate of petition denial. To get a sense of the substantive size of the effect, we can again look at the estimated effect of a move from the 10th to the 90th percentile of automation exposure. This increase in automation exposure is increases the rate of denial by 6 percentage points, roughly a 53% increase over the average commuting zone's denial rate of 16.1 percent.

3.2.3 Discussion

As expected, both the raw number of denied petitions filed per capita and the denial rate are higher in labor markets that have adopted robots as a substitute for labor. These workers seem to believe, at the time of filing, that they have a legitimate case for their job loss being trade related. I interpret this as evidence that these labor markets exhibit a higher overall level of belief that trade is behind job loss.[13]

Although the regressions control for trade shocks from two large sources of trade-related displacement (China and Mexico), they do not control for all possible sources of import competition. Nevertheless, the fact that the number of certified petitions does not increase provides reassurance

[13]Recent work in the use of shift-share instruments has raised concerns about traditional shift-share estimators not adequately controlling for unobserved industry shocks that affect commuting zones with similar industrial specialization. Although explicitly controlling for shocks such as trade is one way to address these concerns, alternative estimation strategies have also been proposed. In Section A.2, I show how my results are robust to alternative estimation strategies proposed by Adão et al. (2019) and Borusyak et al. (2018), which are two prominent papers in this field.

that we have adequately controlled for a commuting zone's import competition. Further, it provides further reassurance for the exclusion restriction. This is because it shows that automation in Europe did not impact TAA activity through its effect on Europe-US trade, since that would show up in an increase in certified petitions.

How does this relate to our larger puzzle about trade and automation *policy*? If workers believe a model of the economy where trade is behind job loss, this belief will show up not just in their beliefs about the appropriate state benefits to apply to, but also in their beliefs about the appropriate national policies to pay attention to. In the next section, I test whether people in these highly exposed commuting zones also have a higher degree of knowledge about trade policy at the national level. If so, this would accord with the idea that automation displaced workers are highly primed on trade policy and the remedies associated with its effects. It would then complement the results presented here, which suggest that these workers not only know more about trade policy at the national level, but that this spills over into how they deal with their own job loss.

3.3 Observable Implication 2: People in Automation Exposed Labor Markets Will Know More About Trade Policy

Measuring Knowledge of Trade Policy: I measure the salience of trade and knowledge of trade policy using the 2006 Cooperative Congressional Election Study (CCES), a nationally representative survey of 36,474 respondents from the United States (Ansolabehere, 2010). A unique feature of the 2006 CCES is that it asks respondents about both their senators' roll call votes on a set of actually proposed legislative bills in the US Congress. The survey data also report the actual vote of both senators on the proposed legislation. This enables me to use the strategy developed by Guisinger (2009) to measure survey respondents' knowledge about their representatives' positions on trade (and non-trade) policy.

The survey first has a short description of the proposal and the arguments for and against it. It then asks how the respondent thinks both senators voted on the issue. Specifically, the 2006 CCES asks questions of the form (this example is for the trade related question):

> *This year Congress also debated a new free trade agreement that reduces barriers to trade between the U.S. and countries in Central America.*
>
> *Some politicians argue that the agreement allows America to better compete in the global economy and would create more stable democracies in Central America. Other politicians argue*

that it helps businesses to move jobs abroad where labor is cheaper and does not protect American producers.

How about Senator 1/2? Do you think he/she voted for or against the trade agreement?

1. For (that is to ratify the trade agreement)

2. Against

3. Don't know

I use answers to questions on four of these issues that were legislatively relevant in 2006, but only one of which was directly relevant to international trade policy. These four issues are: the approval of the Central American Free Trade Agreement (CAFTA, later DR-CAFTA)[14], increasing the Federal Minimum Wage from \$5.15 to \$6.25, the extension of the capital gains tax cut passed in 2001, a measure to ban late-term abortions, a proposed ban on federal funding for stem cell research, a path to citizenship for undocumented immigrants, and setting a deadline for the withdrawal of troops from Iraq. For each of these issues, I code a new binary variable that denotes whether the respondent correctly identified *both* their senators' votes on the issue:

$$Correctly\ Identified\ Senator\ Vote\ _{Senator,Issue} = \begin{cases} 1, & \text{if } RespondentGuess \in \{For, Against\} \\ & \&\ RespondentGuess = SenatorVote \\ 0, & \text{otherwise} \end{cases}$$

$$\begin{aligned} Correctly\ Identified\ Both\ Senators'\ Vote\ _{Issue} = {} & Correctly\ Identified\ Senator\ Vote\ _{1,Issue} \\ & X\ Correctly\ Identified\ Senator\ Vote\ _{2,Issue} \end{aligned}$$

This variable is then equal to one only when the respondent correctly identified both senators' votes on the issue at hand[15].

The values of this variable for the 7 issues I use are presented in the table below:

[14] The Senate had voted 55-45 to approve CAFTA on July 28th 2005.

[15] Alternative configurations were possible; for example I could also create a variable denoting if the respondent correctly identified at-least one senator's vote on an issue. I employ a more strict measure to reduce the chance that a respondent correctly identifies representative positions by random guessing.

Table 5: **Summary Statistics for Survey Variables**:

Proposed Legislation and Respondent Knowledge

Issue	Correctly Identified Both Senators' Vote
Approval of CAFTA	20.2%
Minimum Wage Increase	41.3%
Extension of Capital Gains Tax Cut	10.4%
Late-Term Abortion Measure	31.2%
Stem Cell Research	40.1%
Immigration (Citizenship for Undocumented Immigrants)	29.6%
Troop Withdrawal	43.1%

Note: The percentages are weighted by the survey weight. These reflect answers from 36,421 respondents.

I also code whether the respondent correctly identified the political party of their district's representative in the House of Representatives. Weighted by the survey weight, 69.9% of respondents were able to correctly identify the party of their House representative. In the aggregate, trade seems to be one of the issues where respondents have relatively less information about their political representative's position. However, trade seems more salient than the vote on the extension of the capital gains tax cut, with 20% of respondent being able to correctly identify their representative's position on CAFTA. The relevant question for this analysis, however, is whether this differs by a respondent's exposure to a labor market that is affected by widespread automation.

3.3.1 Design

I first match the respondent's county of residence, provided on the CCES, to their commuting zone of residence. I then test whether respondents who live in more automation exposed commuting zone are more knowledgeable of their representative's policy positions and/or party. I include the representative's party and the other 6 non-trade issue measures as placebos to see if any effects on knowledge about trade policy are reflective of simply a higher degree of knowledge about policy more broadly. The timing of the survey in 2006 is also ideal since it is near the end of the automation exposure time period (2004-2007).

More specifically, I again estimate two-stage least squares regressions of the form:

$$\text{First Stage :}$$

$$Exposure\ to\ Robots : US_{k,c,2004-2007} = \beta_0 + \beta_1 Exposure\ to\ Robots : Europe_{k,c,1993-2007} + \beta\mathbf{X}_{k,c} + \epsilon$$

$$\text{Second Stage :}$$

$$Correctly\ Identified\ Both\ Senators'\ Vote\ _{Issue,k,c} = \alpha_0 + \alpha_1 Exposure\ to\ \widehat{Robots :} US_{k,c,2004-2007} + \beta\mathbf{X}_{k,c} + \gamma$$

Where now the regression is at the individual respondent level, indexed by k. Note however, that the main treatment variables (i.e. the automation exposure variables) vary at the commuting zone level rather than the individual level. In order to correct standard errors for this perfect dependence between observations within commuting zones, I estimate cluster-robust standard errors at the level of treatment; that is, at the commuting zone level (Abadie et al., 2017)[16].

$\mathbf{X}_{k,c}$ is a matrix of individual level control variables from the CCES, including the educational level of the respondent (on a 6-point scale from no high school to post-graduate), the self reported 3 point party identification of the respondent, and a linear term for the age of the respondent. These control for differing levels of knowledge about policy among certain age groups, party supporters, and educational levels. Also included are two sets of fixed effects for what the relevant Senators' votes actually were. This controls for situations where respondents are more likely to correctly identify either a for or against vote, and a for or against vote is systematically related to the automation exposure of a region. I also control, as in the TAA regressions, for the measures of exposure of a respondents commuting zone from import competition from China and Mexico. I also weight all regressions by the survey weight provided in the 2006 CCES to ensure that the conclusions drawn are nationally representative.

3.3.2 Results

Table 6 presents the main estimated effect of automation exposure on the salience of national level trade policy. Column 1 presents the non-instrumental variables OLS regression of correctly identifying CAFTA positions on the exposure of US commuting zones. Columns 3-7 present the main

[16]Aggregating the variables to their commuting zone level means gives substantively the same results, with similar levels of standard errors, and therefore does not change the conclusions reached.

IV second stage estimates, progressively adding control variables.

As column 7 shows, I find a strong positive and statistically significant effect of the exposure of a local labor market to automation and the likelihood that a resident of that labor market will correctly identify both their senators' vote on CAFTA. To put the coefficient in perspective, a change in automation exposure from the 10th percentile to the 90th percentile is estimated to increase the likelihood of knowing political positions on CAFTA by 5.3 percentage points, a roughly 26.2% increase in the likelihood over the average likelihood of 20.2 percent.

Table 6: **Trade Policy is More Salient in Automation-Exposed Areas**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Effect on Trade's Salience	Dependent Variable: *Correctly Identified Both Senators' CAFTA Vote?*						
	OLS	IV: Second Stage					
Exposure to robots: US	0.058***	0.047***	0.047***	0.046***	0.052***	0.091***	0.099***
	(0.011)	(0.014)	(0.014)	(0.014)	(0.014)	(0.016)	(0.019)
Dep Var Mean	0.202	0.202	0.202	0.202	0.202	0.202	0.202
Trade Shock Control?	No	No	No	No	No	No	Yes
Senators' Votes FE?	No	No	No	No	No	Yes	Yes
Education Level FE?	No	No	No	No	Yes	Yes	Yes
Party ID FE?	No	No	No	Yes	Yes	Yes	Yes
Age Control?	No	No	Yes	Yes	Yes	Yes	Yes
# CZs	669	669	669	669	669	669	669
# Respondents	35952	35952	35952	35896	35830	35812	35812
R squared	0.004	0.004	0.008	0.009	0.019	0.045	0.045

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Standard errors are clustered at the commuting zone (CZ) level and reported in parentheses. The unit of analysis is the a survey respondent in the 2006 CCES. All regressions control for the change in the employment to population ratio of a commuting zone from 1990 to 2008. All units are weighted by sampling weights in the CCES.

Although this suggests that automation-related displacement, or the risk of it, causes workers to seek information about trade policy, it could be that economic displacement causes workers to be more attuned to policy in general. I test this formally in table 7, which presents analogues for the 2nd stage IV regression in table 6 column 7 for the other six non-trade issues, as well as for the likelihood that the respondent will correctly identify their representative's party. In columns 1, 2, 4, and 6 of table 7, I fail to find any significant effect of automation exposure on knowledge about senator positions on minimum wage, capital gains, stem cells, or troop withdrawal. The effect sizes are substantively small and too noisy to make robust statistical inferences. In columns 3 and 5, I find that automation exposed labor markets features workers who are *less* likely to know about their senators' positions on the ban on late-term abortions and a path to citizenship for undocumented immigrants. Although these coefficients are not the focus of this paper, they are substantively interesting. For the purposes of the placebo exercise, we can rule out the fear that an increased knowledge of trade is indicative of an increased knowledge of policy more broadly. Finally, I find no effect of automation exposure on being able to correctly identify the political party of a respondents' congressional representative.

Table 7: **Non-Trade Policy is Not More Salient in Automation-Exposed Areas**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	IV: Second Stage						
Effect on Other Issues' Salience	Dependent Variable:						
	Correctly Identified Both Senators' Vote on :						*Correctly Identified*
	Min. Wage	*Capital Gains*	*Abortion*	*Stem Cells*	*Immigration*	*Troop Withdrawl*	*Representative's Party*
Exposure to robots: US	-0.070	-0.006	-0.039**	0.020	-0.121***	0.005	0.006
	(0.059)	(0.017)	(0.016)	(0.025)	(0.019)	(0.017)	(0.010)
Dep Var Mean	0.413	0.104	0.312	0.401	0.296	0.431	0.699
Trade Shock Control?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Senators' Votes FE?	Yes	Yes	Yes	Yes	Yes	Yes	No/NA
Education Level FE?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Party ID FE?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age Control?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# CZs	669	669	669	669	669	669	669
# Respondents	35812	35812	35812	35812	35812	35812	35830
R squared	0.119	0.051	0.187	0.081	0.088	0.061	0.081

Notes: $*p < 0.1$, $**p < 0.05$, $***p < 0.01$. Standard errors are clustered at the commuting zone (CZ) level and reported in parentheses. The unit of analysis is the a survey respondent in the 2006 CCES. All regressions control for the change in the employment to population ratio of a commuting zone from 1990 to 2008. All units are weighted by sampling weights in the CCES.

3.3.3 Discussion

Overall, the results from this section suggest that exposure to automation-related displacement from industrial robots causes workers to be better informed about their representative's views on trade policy, but not on other policies. Another way of stating this is that communities facing exposure to automation may increase the share of their attention devoted to finding out their representative's position on trade issues. In a strategic setting where representatives then want to avoid being punished for their votes in the Senate, trade policy will shift closer to representative's district's preferences.

Of-course, this higher level of knowledge about trade issues could also be a *result* of political representatives messaging to voters about trade issues and their stance on them. In other words, voters might misinterpret their job loss as trade related because their representative tells them so. I am unable to distinguish between these two explanations here. However, in both cases, the national salience of trade policy will be raised, relative to the salience of other policies.

In Appendix Section A.1, I show how the effect observed here only holds in locations where respondents have access to television news coverage of their own state. The effect of automation exposure disappears in areas where respondents receive televised news about a nearby state. This suggests an important media effect in guiding the causal beliefs of displaced communities.

For the purposes of the puzzle of automation, this provides evidence for observable implication 2 of automation exposed communities believing trade to be behind job loss. That is, a higher knowledge of specifically trade policy suggests that extreme policy preferences are at-least party the result of a belief in a model of the economy where trade is to blame for job loss.[17]

4 Conclusion

Correctly attributing the cause of job loss is difficult. In this paper, I show that this difficulty leads automation-displaced communities to erroneously believe that trade is behind the job loss that plagues their communities. This finding contributes to recent efforts to resolve the puzzle of why automation has fewer political opponents than international trade openness. Although popular discussion and the current literature on automation argues that automation gets mis-attributed to trade, the only rigorous evidence we have to establish this is that automation exposed workers exhibit support for a bundle of extreme policy preferences such as protectionism, xenophobia and ethno-nationalism.

It is unclear however, whether these extreme policy preferences are a result of an overall dissatisfaction with the status quo or a genuine belief that trade is behind job loss. In this paper I show that automation exposed communities believe trade to be behind job loss. Given that workers displaced by automation are in a similar skill category as those displaced by trade, they can credibly believe that trade is behind their job loss.

To show that communities displaced by automation actually hold such beliefs, I provide two pieces of evidence: (1) displaced workers in automation-exposed labor markets are more likely to erroneously file a petition for trade adjustment assistance (TAA) than workers in less automation exposed labor markets, and (2) people in these labor markets are more likely to correctly identify

[17] As with the TAA results, these results are robust to alternative estimation strategies proposed by the literature on shift-share analyses, as Appendix Section A.2 shows.

their political representatives' position on trade (but not on other, non-trade issues). Together, these observations imply that these populations are highly primed on trade, which is especially surprising given the volume of research showing that trade is not a more broadly salient issue on which individuals have deep knowledge.

How does the theory offered here differ from explanations that are based on a normative aversion to regulating technology? Rodrik (2018), for example, cites fairness concerns held by the public in developed countries that prevent them from voicing opposition to being displaced by machinery. Although I do not discount this possibility, it is not immediately clear why fairness concerns prevent affected populations to demand compensation from automation-displacement. Future work needs to establish whether specifically workers displaced due to automation also hold fairness norms that prevent them from demanding compensation. The results presented here suggest that the lack of demand for compensation is the result of those most hurt believing that trade was behind their job loss.

How do we square these results with those that find that automation-exposure leads to a shift towards ideologically extreme parties (Im et al., 2019), or towards nativist and protectionist policy preferences (Wu, 2019)? Again, these contributions to the literature establish that displacement due to automation results in a a shift of policy preferences towards the *bundle* of policies that are advocated by more extreme parties. The next step for scholars of the politics of automation is to disentangle this shift. Is the shift due to a general reaction against status-quo policies, or specifically due to a belief that trade openness is behind widespread displacement? Although this paper does not speak to the former, it provides evidence for the later. It suggests that automation leads to political extremism through its effects on workers' causal beliefs. Future research can use specialized tools such as causal mediation analysis to further disentangle this shift in policy preferences.

Another important question raised by this paper is: how do political elites respond to displaced workers mis-interpreting their job loss as trade-related? More pointedly, why do we not see ideological entrepreneurs rise up and attempt to inform automation-displaced workers that they should demand compensation specifically geared towards their type of job loss? Although the reactions of political leaders was not the main concern of this paper, a useful next step would be to measure the reactions of political leaders to automation-related displacement in their district.

For the architects of international market openness, this paper paints a bleak picture. The findings here suggest that compensating the losers from globalization is made difficult not just because of the difficultly of identifying those most hurt, but also by the many petitioners who erroneously think that trade is behind their job loss. Future research can build on this by assessing how a (hypothetical) compensation scheme for automation would impact the demand for protectionist policies or for policies compensating the losers from trade.

Overall, I make progress in resolving the puzzle of why globalization has more political opponents than automation: the political system is unable to accurately classify job loss as automation related because ill-informed workers have an incentive to represent their job loss as trade related.

References

International Federation of Robotics, 2017. URL https://ifr.org/.

Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge. When should you adjust standard errors for clustering? Technical report, National Bureau of Economic Research, 2017.

Daron Acemoglu. Directed technical change. *The Review of Economic Studies*, 69(4):781–809, 2002.

Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from us labor markets. *NBER working paper*, (w23285), 2017.

Daron Acemoglu and Pascual Restrepo. Demographics and automation. Technical report, National Bureau of Economic Research, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: how technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.

Rodrigo Adão, Michal Kolesár, and Eduardo Morales. Shift-Share Designs: Theory and Inference*. *The Quarterly Journal of Economics*, 134(4):1949–2010, 08 2019. ISSN 0033-5533. 10.1093/qje/qjz025. URL https://doi.org/10.1093/qje/qjz025.

Massimo Anelli, Italo Colantone, and Piero Stanig. We were the robots: Automation and voting behavior in western europe. *BAFFI CAREFIN Centre Research Paper*, (2019-115), 2019.

Stephen Ansolabehere. CCES Common Content, 2006, 2010. URL https://doi.org/10.7910/DVN/Q8HC9N.

David Autor, David Dorn, Gordon Hanson, Kaveh Majlesi, et al. *Importing political polarization? The electoral consequences of rising trade exposure*. Number w22637. National Bureau of Economic Research Cambridge, MA, 2016a.

David H Autor, David Dorn, and Gordon H Hanson. Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125(584):621–646, 2015.

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*, 8:205–240, 2016b.

Cameron Ballard-Rosa, Mashail Malik, Stephanie Rickard, and Kenneth Scheve. The economic origins of authoritarian values: evidence from local trade shocks in the united kingdom. In *Annual Meeting of the International Political Economy Society, Nov*, volume 17, page 18, 2017.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. Technical report, National Bureau of Economic Research, 2018.

Italo Colantone and Piero Stanig. The trade origins of economic nationalism: Import competition and voting behavior in western europe. *American Journal of Political Science*, 62(4):936–953, 2018.

Christian Dippel, Robert Gold, and Stephan Heblich. Globalization and its (dis-) content: Trade shocks and voting behavior. Technical report, National Bureau of Economic Research, 2015.

Aaron Eglitis and Milda Seputyte. Robots are coming for europe's jobs. *Bloomberg*, 2018. URL https://www.bloomberg.com/news/articles/2018-11-13/robots-are-coming-for-europe-s-jobs-with-the-east-most-exposed.

James J Feigenbaum and Andrew B Hall. How legislators respond to localized economic shocks: evidence from chinese import competition. *The Journal of Politics*, 77(4):1012–1030, 2015.

Jane Gingrich. Did state responses to automation matter for voters? *Research & Politics*, 6(1): 2053168019832745, 2019.

Alexandra Guisinger. Determining trade policy: Do voters hold politicians accountable? *International Organization*, 63(3):533–557, 2009.

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

Michael J Hicks and Srikant Devaraj. The myth and the reality of manufacturing in america. *Center for Business and Economic Research, Ball State University*, page 6, 2015.

Zhen Jie Im, Nonna Mayer, Bruno Palier, and Jan Rovny. The "losers of automation": A reservoir of votes for the radical right? *Research & Politics*, 6(1):2053168018822395, 2019.

Kirsten Jäger. Eu klems growth and productivity accounts 2017 release - description of methodology and general notes. In *The Conference Board Europe*, 2017. URL http://www.euklems.net/TCB/2018/Metholology_EUKLEMS_2017_revised.pdf.

J Bradford Jensen, Dennis P Quinn, and Stephen Weymouth. Winners and losers in international trade: The effects on us presidential voting. *International Organization*, 71(3):423–457, 2017.

Antti Kaihovaara and Zhen Im. Jobs at risk? task routineness, offshorability and attitudes toward immigration. *Working Paper*, 2018.

Minju Kim and Robert Gulotty. Compensation or consternation?: The role of international competition in domestic trade politics. *Working Paper*, 2018.

Sung Eun Kim and Krzysztof J Pelc. How are workers compensated following trade liberalization? *Working Paper*, 2017.

David A Lake. Open economy politics: A critical review. *The Review of International Organizations*, 4(3):219–244, 2009.

Edward D Mansfield and Diana C Mutz. Support for free trade: Self-interest, sociotropic politics, and out-group anxiety. *International Organization*, 63(3):425–457, 2009.

Daniel Michaels. Foreign robots invade american factory floors. *Wall Street Journal*, 2017.

Daniel J Moskowitz. Local news, information, and the nationalization of us elections. Technical report, Working Paper, https://bit. ly/2W15rjp, 2018.

Mancur Olson. *The logic of collective action*, volume 124. Harvard University Press, 2009.

Sungmin Rho and Michael Tomz. Why don't trade preferences reflect economic self-interest? *International Organization*, 71(S1):S85–S108, 2017.

Dani Rodrik. Populism and the economics of globalization. *Journal of International Business Policy*, 1(1-2):12–33, 2018.

US Department of Labor Employment and Training Administration. Trade adjustment assistance petitions. 2018.

Nichole Wu. Misattributed blame? attitudes towards globalization in the age of automation. *Working Paper*, 2019.

Baobao Zhang. No rage against the machines: Threat of automation does not change policy preferences. *MIT Political Science Department Research Paper*, 2019.

A Appendix

A.1 Media Market Effects

The analysis here splits the result in column 7 of Table 6 into respondents who live in television media markets (Designated Media Markets or DMAs) that are centered in their state ("in-state DMA") or in another nearby state ("out-of-state DMA"). This strategy is based on that developed in Moskowitz (2018), except the measure of media market here is binary.

If the television media in one's state has some role in explaining how automation exposure leads to higher knowledge of trade, then we would expect this effect to be higher for those in 'in-state media markets'. Indeed, as columns 2 and 3 of Table 8 shows, the effect of automation exposure on knowledge of CAFTA holds only within the sample of respondents who live in an in-state DMA and therefore are likely exposed more to their own state's news, relative to those in out-of-state DMAs.

Table 8: **Trade Policy is More Salient in Automation-Exposed Areas**

	(1)	(2)	(3)
Effect on Trade's Salience	Dependent Variable: *Correctly Identified Both Senators' CAFTA Vote?*		
	Full Sample	In-State DMA	Out-of-State DMA DMA
Exposure to robots: US	0.099***	0.105***	-0.092
	(0.019)	(0.021)	(0.078)
Dep Var Mean	0.202	0.209	0.155
Trade Shock Control?	Yes	Yes	Yes
Senators' Votes FE?	Yes	Yes	Yes
Education Level FE?	Yes	Yes	Yes
Party ID FE?	Yes	Yes	Yes
Age Control?	Yes	Yes	Yes
# CZs	669	591	238
# Respondents	35812	31025	4787
R squared	0.045	0.040	0.105

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Standard errors are clustered at the commuting zone (CZ) level and reported in parentheses. The unit of analysis is the a survey respondent in the 2006 CCES. All regressions control for the change in the employment to population ratio of a commuting zone from 1990 to 2008. All units are weighted by sampling weights in the CCES.

A.2 Alternative Shift-Share Estimators

Observable Implication 1: Automation Displaced Workers will Erroneously File for TAA

TAA Results - Robustness to Robustness to Adao, Kolesar, Morales (2019) Estimator

	(1)	(2)	(3)	(4)
	Dependent Variable:			
	# Total Petitions per 10,000 people	*# Denied Petitions per 10,000 people*	*# Certified Petitions per 10,000 people*	*% Petitions Denied*
Exposure to robots: US	1.325	0.839*	0.268	0.131**
	(2.103)	(0.466)	(1.489)	(0.055)
# CZs	722	722	722	722
R squared	0.238	0.080	0.244	0.014

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Standard errors are clustered at the state level and reported in parentheses. The unit of analysis is the a commuting zone. All regressions control for the change in the employment to population ratio of a commuting zone from 1990-2007 and for China and Mexico trade exposure to the commuting zone.

TAA Results - Robustness to Borusyak, Hull and Jaravel (2018) Estimator

	(1)	(2)	(3)
	Dependent Variable:		
	# Denied Petitions per 10,000 people	*# Certified Petitions per 10,000 people*	*% Petitions Denied*
Exposure to robots: US	0.203*	0.490*	0.014*
	(0.105)	(0.297)	(0.008)
# Industries	19	19	19
R squared	0.071	0.030	0.175

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. The unit of analysis is an industry as defined by the IFR.

Observable Implication 2: Automation Exposed Communities Will Know More about Trade Policy

CCES Results - Robustness to Robustness to Adao, Kolesar, Morales (2019) Estimator

	(1)	(2)	(3)	(4)	(5)	(6)	(6)	(7)
			Correctly Identified Both Senators' Vote on:					*Correctly Identified*
	CAFTA	*Min. Wage*	*Capital Gains*	*Abortion*	*Stem Cells*	*Immigration*	*Troop Withdrawal*	*Representative's Party*
Exposure to robots: US	0.135***	-0.146**	-0.008*	0.005	0.028	-0.156***	-0.023	-0.021
	(0.034)	(0.070)	(0.005)	(0.020)	(0.090)	(0.047)	(0.025)	(0.055)
# CZones	669	669	669	669	669	669	669	669
R squared	0.096	0.159	0.119	0.387	0.056	0.127	0.111	0.048

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. The unit of analysis is a commuting zone, and survey variables have been averaged up to the commuting zone level, weighted by survey weight. All regressions control for Senators' votes, China and Mexico trade shocks, the change in the employment to population ration between 1990 and 2007, average respondent age, percent of respondents who identify as Democrats, and percent of respondents that are college educated. All units are weighted by sampling weights in the CCES.

CCES Results - Robustness to Borusyak, Hull and Jaravel (2018) Estimator

	(1)	(2)	(3)	(4)	(5)	(6)	(6)	(7)
			Correctly Identified Both Senators' Vote on:					*Correctly Identified*
	CAFTA	*Min. Wage*	*Capital Gains*	*Abortion*	*Stem Cells*	*Immigration*	*Troop Withdrawal*	*Representative's Party*
Exposure to robots: US	0.046***	-0.029*	0.001	-0.008*	0.008	-0.054***	0.013	0.007
	(0.006)	(0.017)	(0.006)	(0.004)	(0.014)	(0.004)	(0.008)	(0.005)
# Industries	19	19	19	19	19	19	19	19
R squared	0.745	0.061	0.035	0.205	0.025	0.882	0.028	0.146

Notes: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. The unit of analysis is an IFR industry, and all survey variables have been aggregated up to the industry level.