QUANTITATIVE POLITICAL ECONOMY RESEARCH GROUP
Department of Political Economy
King's College London

**Automation and the Future of Work:
How Rhetoric Shapes the Response in Policy Preferences**

QPE Working Paper 2020-12

Karen Jeffrey

July 13, 2020

Automation and the Future of Work:

How Rhetoric Shapes the Response in Policy Preferences

Karen Jeffrey*

3rd July 2020

Abstract: I conduct a survey experiment to test how individuals' preferences for redistributive policies respond to news of their vulnerability to an automation-induced labor market shock. As respondents feel more vulnerable, their preferences for redistributive policies remain constant or *decline*. However, introducing rhetoric that causes respondents to view automation-induced inequality as unfair *increases* preferences for several redistributive policies. The effects are pronounced among more-educated respondents - a group expected to increasingly be affected by automation in future. This suggests that, going forward, rhetoric may become increasingly influential in terms of the political viability of a redistributive policy response to automation.

JEL: D30, D72, H23, H24, O38

Keywords: automation, shocks, inequality, redistribution, rhetoric

*Department of Political Economy, King's College London. karen.jeffrey@kcl.ac.uk. I am grateful to David P. Reichert for many valuable and constructive discussions and to Konstantinos Matakos and Florian Foos for their guidance. I also thank members of the Quantitative Political Economy research group at King's College London, and members of the NORFACE Network's TECHNO group for helpful feedback. I acknowledge grants from King's College London, used to fund the survey experiment. Research approval was granted by the Research Ethics Board at King's College London. The pre-analysis plan and data associated with this article are available at: DOI 10.17605/OSF.IO/389UX. Declarations of interests: None.

1 Introduction

In recent decades, growth in the automation of work has increased inequalities between workers (Aksoy, Ozcan, and Philipp 2020; Van Reenen 2011), and between workers and capital-owners (Autor and Salomons 2018). These trends, coupled with predictions that many more workers will be displaced by automating technologies going forward (Frey and Osborne 2017), have led some politicians to advocate redistributive policies that could decouple growth in automation from growth in inequality.¹ Yet, it is not clear that a majority of voters will support redistributive policies - *even if* those same voters lose out to automation.

Although canonical models of preference formation assumed that voters would support redistribution as long as it was in their economic interests to do so (Meltzer and Richard 1981; Romer 1975), the subsequent literature paints a more complex picture.² One strand of the literature suggests that exposure to certain rhetoric (or persuasive messaging) may cause voters facing a labor market shock to update their policy preferences in a way that is misaligned with their economic interests. For instance, rhetoric that speaks to vulnerable voters’ concerns by “scapegoating” the blame for job insecurity on migrants or imports has been identified as a factor behind the growth in support for nationalist or protectionist policies amongst economically disadvantaged groups - *even though* such policies may not be in voters’ best economic interests (Cochrane and Nevitte 2014; Di Tella and Rodrik 2020).

In this paper, I test the hypothesis that rhetoric of a different nature will suppress support for *redistributive policies* as voters feel increasingly vulnerable to an automation-induced labor market shock. I hypothesize that rhetoric that suggests that taxing “wealth creators” threatens jobs (Berry 2016), or that stigmatizes welfare claimants (Baumberg 2016), will resonate with voters facing job insecurity, and

¹For instance, the US Presidential Candidate, Andrew Yang (2020), has advocated for provision of a universal basic income, while the UK Labour Party (2019) has called for reducing workers’ hours by 20% without reducing pay.

²E.g. Alesina and Giuliano (2011), Alesina, Miano, and Stantcheva (2020), Cavallé and Trump (2015), Georgiadis and Manning (2012), and Rehm, Hacker, and Schlesinger (2012).

will predispose voters to resist or oppose redistributive policies as they feel increasingly vulnerable to automation. I also hypothesize that exposing respondents to rhetoric designed to challenge such narratives, and intended to encourage support for redistribution, will increase support for redistributive policies.

Understanding how policy preferences are likely to respond as more, and different types of, voters are exposed to automation is vital for scholars and policy-makers seeking to understand the economic and political challenges posed by automation. And in an era of “post-truth politics” (Suiter 2016), understanding how exposure to an automation shock interacts with exposure to politicized rhetoric is an important part of the puzzle.

I conduct an online survey experiment involving 2,500 adults resident in the UK. I expect that the results will also be relevant for high-income economies throughout Europe and North America, which face similar labor market disruption from automation (Frey et al. 2016; Frey and Osborne 2017). Respondents are randomly assigned to read an informational treatment before being asked about their preferences for a range of redistributive policies. One third of respondents are randomly assigned to a treatment designed to increase their perceived vulnerability to an automation shock. Another third of respondents are assigned to read the same treatment, accompanied by rhetoric designed to challenge opposition to, and encourage support for, redistributive policies. To obtain a baseline, the remaining respondents are assigned to read a placebo treatment intended to have a neutral effect on the policy preferences under consideration.

I find that as respondents perceive themselves to be more vulnerable to an automation shock, their preferences for redistributive policies *remain constant or decline*. This appears at odds with observational findings by Owen (2020) and Thewissen and Rueda (2019), which suggest that individuals who have been most exposed to automation to date have become more supportive of redistribution in general terms. However, I analyze data from the British Social Attitudes Survey (NatCen Social Research 2019) and find that, although there is a positive association between perceived vulnerability to automation and support for redistribution in general terms, there is no association between perceived vulnerability and support for specific redistributive policies. As such,

the measures used to assess preferences for redistribution may explain this variation in results. One possible explanation for the disparity is that confronting individuals with concrete policy proposals makes the costs associated with redistribution more salient, hence resulting in lower support for redistributive policies.

One notable exception to this trend is the response in preferences for the policy of reducing working time without reducing wages. Reducing working time is the only redistributive policy that respondents become significantly *more* supportive of as they feel more vulnerable to an automation shock - perhaps because reducing working time has the unique (among all policies tested) potential to protect against job loss by redistributing the work that would remain after automation. This may suggest that job-protecting policies would be viewed more favorably by vulnerable workers than policies intended to offer protection from economic hardship via a social safety net. Findings from other studies suggest a similar preference. For instance, Owen (2020) finds that increased exposure to automation is positively associated with preferences for government support to protect jobs in declining industries. Further, Fernández-Albertos et al. (2020) find that priming Spanish survey respondents to consider the potential labor market implications of automation does not alter support for welfare state expansion, but reduces support for policies that would increase job insecurity by accelerating the adoption of automating technologies.

Other experimental studies find that priming respondents to consider the potential labor market effects of automation has no effect on preferences for redistributive policies (Fernández-Albertos et al. 2020; Zhang 2019), leading Zhang (2019) to conclude that the “threat of automation does not change policy preferences”. However, I find that exposing respondents to rhetoric designed to challenge opposition to, and encourage support for, redistribution *prevents* any decline in, or *increases*, preferences for redistributive policies. The results of this study therefore suggest that the way preferences respond to an automation shock is not set in stone, but is influenced by the rhetoric that individuals are exposed to. As such, this study underscores the importance of looking beyond the labor market effects of automation as a determinant of policy preferences, to consider the interaction between exposure to a labor market shock and exposure

to rhetoric. In doing so, this study contributes to the literature that investigates the effects of persuasive communication on political preferences (e.g. Shiller 2017; DellaVigna and Gentzkow 2010; Chong and Druckman 2007; Druckman 2001) - including in the presence of a labor market shock (Cochrane and Nevin 2014; Dal Bó et al. 2018; Di Tella and Rodrik 2020; Wu 2019).

I find that one mechanism via which the rhetoric acts on preferences is by appealing to individuals' sense that an automation-induced exacerbation of inequality would be *unfair*. This aligns with findings from Almås, Cappelen, and Tungodden (2020), which suggest that perceptions of fairness play a key role in individuals' support (or not) for redistribution.

A novel contribution of this study is its future-oriented focus. To date, more-educated workers have tended to be insulated from the negative labor market effects of automation (Kurer and Gallego 2019). However, developments in artificial intelligence are expected to increasingly expose more-educated workers to the negative labor market effects of automation in the coming years (Webb 2019). As such, understanding the response in preferences of this influential group will be key for practitioners seeking to design a politically viable policy response that could limit growth in automation-induced inequality. The experimental treatments used in this study are therefore designed to increase more-educated (as well as less-educated) respondents' perceived vulnerability to an automation shock, thereby enabling an exploration of how policy preferences respond across the education distribution.

I find that policy preferences respond in a heterogeneous manner across education-levels. Respondents who possess a bachelor's degree are the only group to become *less* supportive of redistributive policies overall when informed of their vulnerability to an automation shock. However, the addition of rhetoric significantly increases the group's preferences for redistributive policies. As such, while observational studies find that exposure to automation *to date* has been associated with growth in support for far-right and nationalist political parties (Anelli, Colantone, and Stanig 2018; Dal Bó et al. 2018; Frey, Berger, and Chen 2018; Gingrich 2019; Im et al. 2019), this study suggests that, as more-educated workers are increasingly exposed to automation, a

key political outcome could be greater resistance to redistribution. Nevertheless, the results of this study also suggest that the adoption of carefully designed rhetoric may be capable of counteracting such resistance.

More generally, this study contributes to the literature that uses experimental methods to investigate preferences for various redistributive policies, under different conditions, among different population sub-groups, and where different explanatory mechanisms are at play (e.g. Kuziemko et al. 2015; Fisman et al. 2017; Alesina, Miano, and Stantcheva 2020; Karadja, Mollerstrom, and Seim 2017).

The paper proceeds as follows. Section 2 sets out the data collection procedures, introduces the structure of the survey experiment, and describes the treatments. Section 3 describes the data. Section 4 presents the main results, including the effect of the treatments on respondents’ perceived vulnerability to automation, the response in policy preferences, and an exploration of the mechanisms underpinning these effects. Section 5 presents the results of further analysis of how the treatment effects vary across education-levels. Section 6 concludes.

2 The survey experiment

The instrument used in the survey experiment and the analyses conducted were pre-registered with EGAP before data collection commenced in August 2019.³ The treatments and question-wording for the dependent variables is presented in the Appendix.

2.1 Data collection

Responses were collected during two weeks in August 2019. The survey was hosted on Qualtrics’ online platform and the survey company, Prolific,⁴ was commissioned to recruit a general population sample of 2,500 adults resident in the UK. UK-based members of Prolific’s panel were invited to participate in the study (advertised as

³EGAP ID: 20190807AA.

⁴Peer-reviewed assessments find that Prolific out-performs similar survey platforms (such as MTurk) in terms of data quality and functionality (Palan and Schitter 2018; Peer et al. 2017).

having an estimated completion time of four-to-six minutes) in exchange for 54 pence (equivalent to £5.40 - £8.10 per hour). Prolific used cross-stratification sampling to obtain a sample that was representative of the UK-adult population along interlocking age, gender and ethnicity quotas. Respondents were free to exit the survey at any time and to spend up to twenty minutes answering the questions.

Several features were incorporated into the survey in an effort to ensure the validity of the data collected. Pre-screening questions used for sampling were repeated in the survey to test for continuity. Respondents' IP addresses were used to check that the majority of respondents were located in the UK while completing the survey. Time stamps were included on each page of the survey to identify respondents that spent insufficient time reading the treatment text or answering the dependent variable questions. To encourage respondents to read the treatment text carefully, each treatment was preceded by the statement: "It is very important for the success of our research that you read the following passage carefully before continuing." Pop-up windows encouraging (but not requiring) respondents to answer all questions appeared whenever a respondent tried to move to the next page of the survey without answering all questions on the current page. To aid in the detection of bots and to allow for respondents to flag any issues they encountered whilst taking the survey, an open-text field asking if participants would like to add anything about their responses was included on the final page.

2.2 The structure of the survey experiment

The survey experiment consisted of four sections. First, respondents were asked to answer questions about their socio-demographic characteristics and political leanings. Second, respondents were randomly assigned to read one of three informational treatments: (i) the "Automation Shock" treatment, designed to increase perceived vulnerability to an automation shock; (ii) the "Automation Shock with Rhetoric" treatment, designed to increase perceived vulnerability to an automation shock and to challenge opposition to, and encourage support for, redistribution; or (iii) a placebo treatment designed to have a neutral effect on the policy preferences under consideration. Re-

spondents were then asked about their preferences for a range of redistributive policies. In the final section, respondents were asked about their expectations about the likely labor market consequences of automation, as well as other beliefs that might influence their policy preferences.

2.3 The treatments

In the first instance, the treatments were designed to increase respondents' perceived vulnerability to an automation shock and to encourage respondents to believe that a range of workers - including highly-educated workers - might be negatively affected. Although the survey experiment setting is likely to artificially increase attentiveness, it seems reasonable to expect that exposure to a brief, online passage, warning of a threat that is "10 to 15 years" in the future, would produce a more muted effect on respondents' policy preferences than first-hand experience of a shock, or repeated exposure to media coverage of a shock. In an effort to strengthen the treatment effect, a pessimistic account of what automation might mean for workers is therefore presented.

The "Automation Shock" treatment states that "Experts predict that robots and artificial intelligence will destroy 1 in 3 jobs in the UK in the next 10 -15 years", echoing predictions made by Frey and Osborne (2014). The text then lists examples of the types of jobs expected to be automated, ranging from shop assistants to medical specialists, and emphasizes that "For the first time, people with a university education are also thought to be at risk." The text goes on to mention the possible implications of such a shock, stating that "with so many jobs expected to be lost going forward many workers could be forced to accept lower wages or face long periods of unemployment." The treatment ends with a general statement about the need for policy intervention, without mentioning any particular policy direction: "We need to do something, but it's not clear what. None of the major political parties have set out a concrete plan of action to address the issue yet."

The "Automation Shock with Rhetoric" treatment includes the same text as the "Automation Shock" treatment, but replaces the general statement about the need for policy intervention with rhetoric designed to challenge resistance to, and encourage

support for, a redistributive policy response. In an effort to prevent possible resistance to redistribution arising from the belief that taxing “wealth creators” will threaten jobs, the rhetoric re-frames “wealth creators” as benefiting from automation at the expense of workers, stating: “Some say that if corporations replace employees with machines, billionaire bosses and shareholders will become even richer while ordinary workers lose their jobs.” By emphasizing the discrepancy between “billionaire bosses” and “ordinary workers”, the treatment also aims to appeal to respondents’ sense of fairness - a consideration that Almås, Cappelen, and Tungodden (2020) identify as being fundamental to individuals’ acceptance, or not, of inequality. The rhetoric ends by presenting a positive vision of welfare, designed to counteract possible resistance to redistribution stemming from the belief that welfare claimants are undeserving or dishonest, stating: “But if we can find a way to share the wealth that’s created, we could all work less and spend more time doing the things we enjoy.” The emphasis on sharing the wealth created by automation widely is informed by investigations by Korpi and Palme (1998), which suggest that narrowly targeting benefits appears to have a stigmatizing effect that reduces support for welfare provision.

The placebo text is used to obtain a baseline and is therefore designed to have a neutral effect on respondents’ expectations about automation and on the policy preferences under consideration. It discusses changes in the types of jobs that people have been employed to do over the years without making reference to automation. For instance, the treatment mentions the decline in numbers of milkmen and TV repairers and growth in the number of people employed as administrators and care workers.

In an effort to increase engagement, at the end of each treatment respondents are asked to indicate whether they personally know anyone employed in each of the occupations mentioned in the excerpt.⁵

⁵To avoid introducing bias, the order that options are listed in is randomized for each respondent.

3 Data

3.1 Attrition, data cleaning and missing responses

Attrition was low; just 12 of the 2,500 respondents (0.5%) exited the survey without marking their responses as complete. To minimize possible noise resulting from random answer selection by very quick respondents (Greszki, Meyer, and Schoen 2015), I remove responses where less than half the median time was spent on the treatment page or on answering the dependent variable questions. I also remove responses where none of the dependent variable questions were answered. In total, 204 responses were removed, representing 8% of the sample.⁶ The analysis presented in this paper is based on the remaining 92% of responses.⁷ I impute missing covariate data by taking the mean of non-missing values for the variable in question. Where a subject’s response is missing for one of the dependent variables, I exclude the subject from the analysis of that variable.

3.2 Summary statistics

I calculate and apply post-stratification weights so that the sample closely matches the demographic composition of the UK population.⁸ Table 1 presents weighted summary statistics for this study’s sample alongside summary statistics from the UK Census (ONS 2011) and the nationally representative samples of adults surveyed in the European Social Survey (ESS 2016) and the British Social Attitudes Survey (NatCen Social

⁶Removed responses were fairly evenly distributed across treatment groups: 33% from the “Automation Shock” group, 35% from the “Automation Shock with Rhetoric” group, 32% from the placebo group.

⁷The results presented throughout this paper remain substantively unchanged when performed using the unrestricted sample.

⁸I create weights based on the representative sample collected during the 2019 British Social Attitudes Survey (NatCen Social Research 2019) with respect to 72 cells based on: gender (2) x age brackets (2) x white versus non-white (2) x education level (3) x voting in the previous general election (3). The range of the calculated weights is relatively small (ranging from 0.4 to 2.0); as such, I do not trim the weights.

Research 2019).

3.3 Randomization

To minimize random variation across the three treatment groups I use Qualtrics’ Web Service feature to connect the survey to a cloud-based algorithm that performs sequential covariate-adaptive block randomization (following Cavaille n.d.). The algorithm ("**seqblock**", developed by Moore and Schnakenberg (2016)) assigns each respondent to one of the treatment groups based on a measure of multivariate distance from all previously-assigned respondents, calculated on the fly using respondents’ answers to the socio-demographic and political-leanings questions asked at the beginning of the survey. Respondents are therefore randomly assigned (with equal probability) to one of the three treatment groups whilst blocking on age, gender, ethnicity, education-level, political-leaning and employment status.

Table 2 presents summary statistics stratified by treatment group and indicates statistically significant differences in the composition of each group compared to the baseline (Placebo group). Of the 34 variables tested (17 for each treatment group), only one yields a statistically significant result from a chi-squared test of independence at $p < 0.05$, consistent with successful randomization. Tests for joint orthogonality also show that the socio-demographic composition of the “Automation Shock” treatment group and the “Automation Shock with Rhetoric” treatment group do not vary significantly from the “Placebo” group ($p = 0.53$ and $p = 0.39$, respectively).

4 Results

I present three sets of results. First, I present the effects of the treatments on individuals’ perceived vulnerability to an automation shock. Second, I present the effects of the treatments on individuals’ policy preferences. Third, I present analysis of the possible mechanisms via which the treatments act on individuals’ policy preferences.

Table 1: Summary statistics and comparison to other surveys

	This study's sample (2019)	UK Census (2011)	European Social Survey (2016)	British Social Attitudes Survey (2019)
<i>Socio-demographics (%)</i>				
Male	48.8	48.6	48.9	48.8
Age 18 - 34	28.5	29.1	28.8	28.5
Age 35 - 54	35.7	35.3	34.1	34.0
Age 55+	35.8	35.6	37.1	37.5
White	87.1	87.3		87.1
Asian	6.4	7.2		7.0
Black	3.3	3.1		4.1
Other ethnicity	3.2	2.4		1.8
Employed (full or part)	61.1		57.7	56.9
Unemployed	5.0		3.6	4.5
Student	5.3		5.6	5.0
Has university degree	27.9		29.8	27.9
Conservative voter	28.9		28.7	28.9
Labour voter	32.1		26.3	32.1
Observations	2, 296	41.7M	1, 959	3, 988

Notes: Column 1 displays summary statistics from this study's weighted regression sample (the 92% of the sample remaining after cleaning). Column 2 displays summary statistics from the UK Census (ONS 2011), excluding individuals under 18 years of age. Column 3 displays (weighted) statistics from the 8th European Social Survey (ESS 2016), excluding individuals under 18 years of age or not resident in the UK. Column 4 displays (weighted) statistics from the 35th British Social Attitudes Survey (NatCen Social Research 2019).

Table 2: Balance across treatment groups

	Placebo	Automation Shock	Automation Shock with Rhetoric
<i>Socio-demographics (%)</i>			
Male	47.7 (0.5)	49.0 (0.5)	49.6 (0.5)
Age 18 - 34	28.5 (0.5)	28.3 (0.5)	28.7 (0.5)
Age 35 - 54	35.8 (0.5)	36.7 (0.5)	34.6 (0.5)
Age 55+	35.8 (0.5)	35.0 (0.5)	36.7 (0.5)
White	88.8 (0.3)	86.7 (0.3)	85.6 (0.4)
Asian	6.1 (0.2)	6.3 (0.2)	6.9 (0.3)
Black	2.0 (0.1)	4.0** (0.2)	3.9* (0.2)
Other ethnicity	3.2 (0.2)	3.0 (0.2)	3.6 (0.2)
Low income (<£20k)	25.2 (0.4)	25.0 (0.4)	26.4 (0.4)
Medium income (£20k - £60k)	61.2 (0.5)	58.9 (0.5)	58.4 (0.5)
High income (>£60k)	13.6 (0.3)	16.1 (0.4)	15.2 (0.4)
Employed (full or part)	59.9 (0.5)	60.1 (0.5)	63.2 (0.5)
Unemployed	5.5 (0.2)	4.4 (0.2)	5.1 (0.2)
Student	5.4 (0.2)	6.2 (0.2)	4.3 (0.2)
Has university degree	27.8 (0.4)	26.6 (0.4)	29.4 (0.5)
Conservative voter	30.1 (0.5)	28.1 (0.4)	28.4 (0.5)
Labour voter	33.7 (0.5)	32.4 (0.5)	30.1* (0.5)
Observations	764	770	762

Notes: Mean variable values for each treatment group are presented in the first row. Standard deviations are presented in parentheses. Asterisks indicate significant p-values from chi-square tests of independence between treatment assignment and observed covariates. All statistics presented in this table are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning). ***p<0.01, **p<0.05, *p<0.1.

4.1 Effect of treatments on perceived vulnerability to an automation shock

I test the effect of the treatments on two measures of perceived vulnerability to an automation shock: perceived societal vulnerability and perceived personal vulnerability. Perceived *societal* vulnerability is measured by asking respondents to what extent they agree that: “More jobs will be lost than gained as a result of automation by robots or artificial intelligence in the next 10 - 15 years.” Perceived *personal* vulnerability is measured by asking respondents to what extent they agree with the statement: “It is likely that you or a member of your family will lose your job to robots or artificial intelligence in the next 10 – 15 years.” Agreement with each statement is coded on a seven-point scale standardized from zero (strongly disagree) to 100 (strongly agree). Respondents were asked about their perceived vulnerability *after* being asked about their policy preferences to avoid priming members of the Placebo group to think about automation while responding to the policy preference questions.

Baseline agreement with the societal vulnerability measure is considerably higher than baseline agreement with the personal vulnerability measure (72 compared to 45 on a scale from zero to 100). That respondents tend to view automation as something that threatens other people, but not themselves or their relatives, echoes existing survey results (for instance, Ipsos (2018) and Phillips et al. (2018)).

Figure 1 presents the estimated marginal effects of the treatments on each measure of perceived vulnerability, relative to the baseline. The plotted estimates come from linear regressions,⁹ of the form shown in Equation 1.

$$\text{Vuln}_i = \beta_0 + \beta_1 \text{AutoShock}_i + \beta_2 \text{AutoShockRhetoric}_i + \beta_3 \mathbf{Z}_i + \epsilon_i \quad (1)$$

Vuln_i is individual i ’s perceived vulnerability to an automation shock, AutoShock_i and $\text{AutoShockRhetoric}_i$ are binary variables signaling i ’s exposure (1) or not (0) to each treatment, and $\mathbf{Z}_i$ is a vector of controls containing all socio-demographic and political variables listed in Table 2. The inclusion of control variables does not change the results substantively, consistent with successful randomization. Therefore, to conserve

⁹The results remain substantively unchanged when estimated using ordered logistic regression.

space and minimize noise, I only present results from regressions that include control variables.

Figure 1: Marginal effect of treatments on perceived vulnerability to automation

Notes: Estimates from: $Vuln_i = \beta_0 + \beta_1 AutoShock_i + \beta_2 AutoShockRhetoric_i + \beta_3 Z_i + \epsilon_i$. 95% confidence intervals are shown. The significance of the marginal effect of the “Automation Shock with Rhetoric” treatment relative to the “Automation Shock” treatment is shown outside the square bracket on the right-hand side. The outcome variables are coded on a seven-point scale standardized from zero (strongly disagree) to 100 (strongly agree) and signal agreement with the statements: (i) More jobs will be lost than gained as a result of automation by robots or artificial intelligence in the next 10 - 15 years, and (ii) It is likely that you or a member of your family will lose your job to robots or artificial intelligence in the next 10 - 15 years. Binary variables are used to signal respondents’ exposure (1) or not (0) to each treatment. Estimates should be interpreted as the change in percentage points relative to the base category. Controls are included for all variables listed in Table 2. All statistics presented in this figure are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

As shown in Figure 1, the treatments have no significant effect (at $p < 0.05$) on individuals’ perceived *societal vulnerability* to automation, relative to the baseline. This null-effect likely reflects the high baseline agreement with the measure of societal vulnerability. However, the “Automation Shock” and the “Automation Shock with Rhetoric” treatments cause large, statistically significant increases in individuals’ perceived *personal vulnerability* - by 8% (or 5.2 percentage points, $SE = 1.6$, $p < 0.01$) and 23% (or 10.5 percentage points, $SE = 1.6$, $p < 0.01$) relative to the baseline, respectively. Descriptive statistics presented in Figure 2 show that a subset of treated respondents in all age, education, and income groups perceive society, or themselves, to be vulnerable to automation; however, perceived vulnerability is lower among the most-educated and

highest-earning respondents. Taken together, these findings suggest that the treatments achieve the intended effect of increasing respondents’ perceived vulnerability to an automation shock, and that individuals across the education distribution perceive themselves to be vulnerable.¹⁰

Before proceeding, to rule out the possibility of interference with the theorized causal mechanism, I test whether either treatment diminishes respondents’ trust in government (a factor that has been shown to diminish support for redistributive policies (Dal Bó et al. 2018; Kuziemko et al. 2015; Van Oorschot 2006)). Reassuringly, I find that neither treatment significantly influences respondents’ trust in government.¹¹

4.2 Effect of the treatments on policy preferences

Figure 3 presents the estimated effects of each treatment on respondents’ agreement that, going forward, the government should implement each of seven redistributive policies. In light of observational findings that exposure to automation is associated with increased support for nationalist political parties (Anelli, Colantone, and Stanig 2018; Frey, Berger, and Chen 2018; Dal Bó et al. 2018; Im et al. 2019; Gingrich 2019), I also test the effect of treatments on one nationalist policy (restricting immigration).

Policy preferences are coded on a seven-point scale standardized from zero (strongly disagree with the policy being implemented) to 100 (strongly agree with the policy being implemented). Results are estimated using linear regressions¹² as shown in Equation 2.

$$\text{PolicyPref}_i = \beta_0 + \beta_1 \text{AutoShock}_i + \beta_2 \text{AutoShockRhetoric}_i + \beta_3 \mathbf{Z}_i + \epsilon_i \quad (2)$$

PolicyPref_i is the extent to which individual i agrees that the government should implement the policy going forward; AutoShock_i and $\text{AutoShockRhetoric}_i$ are binary vari-

¹⁰Figure 7, Panel (d) provides further evidence of these effects.

¹¹Trust in government is measured as agreement with the statement: “In general, the major political parties in the UK can be trusted to deliver on the promises they make”. The estimated effect of the “Automation Shock” and the “Automation Shock with Rhetoric” treatments on trust in government are -1.45 percentage points (SE=1.23) and -2.05 percentage points (SE=1.17), respectively.

¹²The results remain substantively unchanged when estimated using ordered logistic regression.

Figure 2: Perceived vulnerability to automation, by demographic group

Notes: Plots (a) - (d) show the percentage of respondents who received either the "Automation Shock" or the "Automation Shock with Rhetoric" treatment, who agree (or not) with the statement "More jobs will be lost than gained as a result of automation by robots or artificial intelligence in the next 10 - 15 years." Plots (e) - (h) show the percentage that agree (or not) with the statement "It is likely that you or a member of your family will lose your job to robots or artificial intelligence in the next 10 - 15 years." Plots (b) and (f) show agreement within each age band. Plots (c) and (g) show agreement within each education band. Plots (d) and (h) show agreement within each income level. All statistics presented are based on the weighted regression sample (the 92% of the sample remaining after cleaning).

ables signaling i 's exposure (1) or not (0) to the “Automation Shock” treatment and the “Automation Shock with Rhetoric” treatment; and $\mathbf{Z}_i$ is a vector of controls containing all socio-demographic and political variables listed in Table 2.

To minimize the false discovery rate, I adjust p-values for multiple comparisons, following Benjamini and Hochberg (1995).¹³ The estimated coefficients on the control variables reveal patterns that are consistent with existing investigations. Support for redistributive policies decreases as income increases. Conservative Party voters are less supportive of redistributive policies, and are more supportive of restricting immigration, while Labour Party voters are more supportive of redistributive policies.

I test preferences for redistributive policies that would redistribute to unemployed or low-earning individuals, as well as policies that would redistribute away from high-earning individuals and corporations. I also include two less commonly adopted redistributive policies that have been proposed in response to automation: providing a universal basic income, and reducing working time without reducing pay. Baseline support for the redistributive policies varies considerably. The Placebo group favors policies that would redistribute away from the rich – whether via income tax or corporation tax (mean agreement is 73 and 77, respectively, on a scale from zero to 100) as well as policies to redistribute in favor of low-earners, such as increasing the minimum wage and increasing income support (mean agreement is 79 and 67, respectively). Support for a universal basic income is lower (mean agreement is 60) and support for reducing working time and increasing unemployment benefit is lowest of all (mean agreement is 51 and 50, respectively).

Effect of the “Automation Shock” treatment on preferences for redistributive policies

Looking first at the effect of the “Automation Shock” treatment on respondents’ preferences for redistributive policies, Figure 3 shows that, despite increasing respondents’ perceived vulnerability, the “Automation Shock” treatment causes a significant decline

¹³Estimates with unadjusted p-values, including coefficients for all control variables are detailed in Appendix Table 3

in preferences for two of the redistributive policies: increasing income support for low-earners and increasing the minimum wage. For robustness, I test for the possibility that the original Placebo treatment inflated baseline preferences for increasing income support and increasing the minimum wage by mentioning low-paying occupations. I re-administer the survey to a sample of 300 respondents, who are exposed to a thematically unrelated Control treatment (on the medicinal uses of honey). I find that preferences for increasing income support and increasing the minimum wage do not differ significantly between the original Placebo group and the follow-up Control group.¹⁴

The treatment has no significant effect on preferences for the other redistributive policies (at $p < 0.05$) - with one notable exception. The “Automation Shock” treatment *increases* support for the policy of reducing working time with no reduction in pay, by 6.0 percentage points (SE=1.6, adjusted p -value <0.01). This may reflect the policy’s unique (among all policies tested) potential to protect against job loss by redistributing the work that remains after automation. In other words, this suggests that as respondents feel more vulnerable to an automation shock, they favor job-protecting policies over welfare provision. Findings from Fernández-Albertos et al. (2020) and Owen (2020)) appear to suggest a similar preference.

The decline in support for, or null-effects on, preferences for all other redistributive policies in response to the “Automation Shock” treatment appears somewhat at odds with findings that suggest that exposure to automation is positively associated with support for redistribution in general terms¹⁵ (Owen 2020; Thewissen and Rueda 2019). However, as shown in Figure 4, analysis of data from the British Social Attitudes Survey (BSAS) (NatCen Social Research 2019) reveals that BSAS respondents who are more concerned about their vulnerability to automation¹⁶ *are* more supportive of redistribu-

¹⁴The full text of the new Control is included in the Appendix, and the comparison of policy preferences are summary statistics for both groups is presented in Appendix Table 4.

¹⁵Measured as agreement with the statement: “The government should take measures to reduce differences in income levels” (Thewissen and Rueda 2019) or “Do you think it is the government’s responsibility to reduce income differences between the rich and poor” (Owen 2020).

¹⁶BSAS respondents are asked: “How worried, if at all, are you that in the next 10 years, your job might be done instead by machines or computer programmes?” on a four-point scale from “Not at all worried” to “Very worried”.

Figure 3: Marginal effect of treatments on policy preferences

Notes: Estimates from: $\text{PolicyPref}_i = \beta_0 + \beta_1 \text{AutoShock}_i + \beta_2 \text{AutoShockRhetoric}_i + \beta_3 \mathbf{Z}_i + \epsilon_i$. 95% confidence intervals are shown. The marginal effect of the “Automation Shock with Rhetoric” treatment relative to the “Automation Shock” treatment is shown outside the square bracket on the right-hand side. The outcome variables are coded on a seven-point scale standardized from zero (strongly disagree) to 100 (strongly agree) and signal agreement with the statements: “Going forward, the government should. . . (i) increase benefits paid to people with low incomes; (ii) increase the minimum wage that employers can legally pay their staff; (iii) increase taxes on incomes of over £150,000; (iv) increase taxes on corporations’ profits; (v) increase benefits paid to people who are unemployed; (vi) pay every citizen enough to cover the cost of basic necessities, regardless of whether they work or not; (vii) make it possible for workers to reduce their working hours by 20% with no reduction in pay”; (viii) increase restrictions on the number of immigrants allowed to work in the UK”. Binary variables are used to signal respondents’ exposure (1) or not (0) to each treatment. Estimates should be interpreted as the change in percentage points relative to the base category. Benjamini and Hochberg’s method adjusts p-values for multiple comparisons. All variables listed in Table 2 are included as controls. All statistics presented in this figure are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

tion in general terms (in-line with Owen (2020) and Thewissen and Rueda (2019)), but *are not* significantly more likely to support specific redistributive policies, including: increasing welfare, increasing the minimum wage, and increasing unemployment benefit. One possible explanation for this discrepancy is that confronting individuals with concrete policy proposals makes the costs associated with redistribution more salient, resulting in weaker or negative preferences for redistributive policies.

Figure 4: Perceived future vulnerability to automation and preferences for redistribution and redistributive policies

Notes: Estimated coefficients from OLS regressions on the British Social Attitudes Survey (NatCen Social Research 2019). Estimates from the model: $\text{PolicyPref}_i = \beta_0 + \beta_1 \text{PerceivedVuln}_i + \beta_3 \mathbf{Z}_i + \epsilon_i$, where PolicyPref_i represents responses to: (i) “It is the responsibility of the government to reduce the differences in income between people with high incomes and those with low incomes.”; (ii) “Government should redistribute income from the better-off to less well-off”; (iii) “The government should spend more money on welfare benefits for the poor”; (iv) “Do you think the minimum wage should be increased, reduced or kept at the same level it is now?”; (v) “Would you like to see more or less government spending than now on benefits for unemployed people?” PerceivedVuln_i represents responses to the question: “How worried, if at all, are you that in the next 10 years, your job might be done instead by machines or computer programmes?” $\mathbf{Z}_i$ controls for sex, age, ethnicity, employment status, education, voting in the previous general election, and income. Responses are re-coded so that estimates represent the change (in percentage points) in preferences for each policy option for a one unit increase in perceived future vulnerability to automation. Data are weighted using the supplied weights: WtFactor . $N = 533 - 850$. 95% confidence intervals are shown.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Effect of the “Automation Shock with Rhetoric” treatment on preferences for redistributive policies

Turning next to the effect of the “Automation Shock with Rhetoric” treatment, Figure 3 shows that presenting news of an automation shock alongside rhetoric designed to encourage support for redistribution *prevents* the decline in preferences for increasing income support and increasing the minimum wage and causes significant *increases* in preferences for increasing corporation tax, increasing unemployment benefit, providing a universal basic income, and reducing working time without reducing wages. The only remaining redistributive policy that does not see a significant increase in support (at $p < 0.05$) is the policy of providing a universal basic income.

In summary, while the “Automation Shock” treatment reduces or has no effect on preferences for some redistributive policies, the “Automation Shock with Rhetoric” treatment prevents any reduction in preferences and significantly increases support for all-but-one of the remaining redistributive policies.

Effect of the treatments on preferences for restricting immigration

Figure 3 shows that neither treatment has a significant effect on preferences for restricting immigration. However, disaggregating the results by income-level reveals that the “Automation Shock” treatment causes individuals in households in the bottom-half of the income distribution to become significantly *more* supportive of restricting immigration (by 4.2 percentage points, $SE = 2.1$, $p < 0.05$). This is consistent with the results of existing observational studies, which find that workers who have been exposed to automation *to date* (who tend to be situated around the lower-half of the income distribution (Kurer and Gallego 2019), as theorized by Autor, Levy, and Murnane (2003)) are more likely to support nationalist and far-right political parties. The “Automation Shock with Rhetoric” treatment does not cause any significant change in preferences for restricting immigration at the aggregate or disaggregated level.

Normalizing effect sizes with respect to baseline preferences

To account for variation in preferences for each of the policies at the baseline, I normalize the treatments’ effect sizes with respect to baseline preferences by calculating persuasion rates. The persuasion rate, developed by DellaVigna and Kaplan (2007) and simplified by Di Tella and Rodrik (2020), represents the percentage change in the share of respondents who did not agree with a policy at the baseline, who are persuaded to agree with the policy following exposure to treatment, as shown in Equation 3.

$$\text{PersuasionRate} = 100 * \frac{y_T - y_B}{1 - y_B} \quad (3)$$

y_T is the share of respondents in a treatment group who agree¹⁷ that, going forward, the government should implement the policy in question. y_B is the share of respondents at the baseline (Placebo group) who agree that government should implement the policy going forward. If baseline support for a policy (y_B) is high, the denominator in Equation 3 will be smaller ($1 - y_B$), such that the persuasion rate will tend to be larger if more respondents agree with the policy at the baseline. Vice versa, if baseline support for a policy (y_B) is low, the denominator in Equation 3 will be larger ($1 - y_B$), such that the persuasion rate will tend to be smaller if fewer individuals agree with the policy at the baseline. As such, the persuasion rate incorporates an element of the *potential* for the treatments to move preferences for each policy, making the relative effect sizes more comparable in light of variation in preferences at the baseline.

Figure 5 presents the persuasion rates for each treatment. While preferences for each policy generally move in the same direction as indicated in Figure 3, the relative magnitudes of the effect sizes differ. For instance, following exposure to the “Automation Shock” treatment, the decline in support for increasing the minimum wage is considerably larger than the decline in support for increasing income support when the effects are normalized with respect to baseline preferences (-41% compared to -18%, respectively). Further, the increase in support for increasing corporation tax appears larger

¹⁷Respondents who indicate that they “Strongly agree”, “Agree”, or “Somewhat agree” with the policy, and excluding respondents who indicate that they “Neither agree nor disagree”, “Somewhat disagree”, “Disagree” or “Strongly disagree”.

than the increase in support for reducing working time when normalized with respect to baseline preferences (30% compared to 24%, respectively).

To put the scale of these effects in context, persuasion rates above 20% suggest a relatively large persuasive effect when bench-marked against a range of studies on the effects of consumer and political marketing considered by DellaVigna and Gentzkow (2010). According to this guide, the “Automation Shock” treatment causes a relatively large decline in preferences for increasing income support and increasing the minimum wage, and a relatively large increase in support for reducing working time. The “Automation Shock with Rhetoric” treatment causes a large increase in support for increasing corporation tax, increasing the top rate of income tax and reducing working time.

4.3 Why does the addition of rhetoric increase support for redistributive policies?

Next, I investigate *why* the “Automation Shock with Rhetoric” treatment increases preferences for redistributive policies. Figure 6 presents the marginal effect of the “Automation Shock with Rhetoric” treatment, relative to the effect of the “Automation Shock” treatment on three beliefs that the rhetoric was designed to speak to (discussed in section 2.1): (i) the belief that increasing taxes on “wealth creators” leads to job losses;¹⁸ (ii) the belief that growth in automation-induced inequality would be unfair;¹⁹ and (iii) negative perceptions of welfare claimants.²⁰

The effect of the “Automation Shock with Rhetoric” treatment is not significantly different from the “Automation Shock” treatment in terms of the belief that increasing taxes on “wealth creators” leads to job losses, or in terms of negative perceptions of welfare claimants. However, respondents’ belief that growth in automation-induced

¹⁸Measured as (dis)agreement that: “Increasing taxes on corporations usually leads to job losses”.

¹⁹Measured as (dis)agreement that: “It would be unfair if the gap between those with high and low incomes were to increase as a result of workers being replaced by robots or artificial intelligence”.

²⁰Measured as (dis)agreement that: “Benefits for unemployed people are too high and discourage people from working”.

Figure 5: Persuasion rates

Notes: Each bar shows the percentage change in the share of respondents who did not agree with a policy at the baseline, who are persuaded to agree with the policy following exposure to treatment. Persuasion rates are calculated: $\text{PersuasionRate} = 100 * \frac{y_T - y_B}{1 - y_B}$, where y_T is the share of respondents in the treatment group who “Strongly agree”, “Agree”, or “Somewhat agree” with the policy being implemented, and y_B is the share of respondents who “Strongly agree”, “Agree”, or “Somewhat agree” with the policy being implemented at the baseline (Placebo group). Mean baseline support for each policy is shown in brackets under the policy name. All statistics presented in this figure are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning).

inequality would be unfair increases by 3.5 percentage points ($SE = 1.1$, $p < 0.01$) relative to the effect of the “Automation Shock” treatment. Within the Placebo group, a one percentage point increase in the belief that “It would be unfair if automation increased inequality” is associated with a 30 - 36 percentage point increase in support for each of the redistributive policies.

Taken together, these findings suggest that one mechanism via which the addition of rhetoric increases respondents’ preferences for redistributive policies is by appealing to their sense that growth in automation-induced inequality would be unfair.

Figure 6: Marginal effect of rhetoric on beliefs

Notes: Estimates from: $\text{Belief}_i = \beta_0 + \beta_1 \text{AutoShockRhetoric}_i + \beta_2 \text{Placebo}_i + \beta_3 \mathbf{Z}_i + \epsilon_i$. 95% confidence intervals are shown. The outcome variables are coded on a seven-point scale standardized from zero (strongly disagree) to 100 (strongly agree) and signal agreement with the statements: (i) “Increasing taxes on corporations usually leads to job losses”; (ii) “It would be unfair if the gap between those with high and low incomes were to increase as a result of workers being replaced by robots or artificial intelligence”; and (iii) “Benefits for unemployed people are too high and discourage people from working”. Binary variables are used to signal respondents’ exposure (1) or not (0) to each treatment. Estimates should be interpreted as the marginal effect of the “Automation Shock with Rhetoric” treatment relative to the “Automation Shock” treatment, in percentage points. Benjamini and Hochberg’s method adjusts p-values for multiple comparisons. All variables listed in Table 2 are included as controls. All statistics presented in this figure are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5 Heterogeneity in preferences by education-level

In this final section, I test for heterogeneous effects of the treatments on policy preferences across education-levels. I group respondents according to their highest level of education (from: (i) some high-school or less; to (ii) further high-school or some college; (iii) bachelor’s degree; or (iv) postgraduate, professional degree or doctorate), before estimating the effect of the treatments on the policy preferences of each group.

The effect of the treatments on preferences for restricting immigration is relatively homogeneous across the four education-levels; however, the effect on preferences for the seven redistributive policies is more heterogeneous (illustrated in Appendix Figure 8). To obtain a generalized picture of how the treatments influence preferences for redistributive policies, I use principal component analysis to construct an index that aggregates preferences for the seven redistributive policies.²¹ Figure 7, Panel (a) presents the estimated effects of the treatments on preferences for redistributive policies, aggregated according to the index.²²

As shown in Figure 7, Panel (a), the “Automation Shock with Rhetoric” treatment has no additional effect (over and above the “Automation Shock” treatment) on the least- and most-educated respondents’ preferences for redistributive policies. However, the “Automation Shock with Rhetoric” treatment significantly increases support for redistributive policies among the middle-two education-levels. In particular, for respondents with a bachelor’s degree, the addition of rhetoric reverses the sign of the coefficient from negative to positive.

I explore whether the heterogeneous response in policy preferences across education-levels is explained by variation in concerns about automation-induced inequality, or in perceived societal or personal vulnerability to automation. However, the results presented in Figure 7, Panels (b), (c) and (d) suggest that these factors do not explain the heterogeneous response in preferences across education-levels.

²¹Factor loadings for the first principal component range from -0.41 to -0.35.

²²The results of this analysis remain substantively unchanged when the weighted mean of preferences for the seven redistributive policies is used in place of the index.

Figure 7: Estimated effect of treatments on preferences, by education-level

Notes: Estimates from: $Y_i = \beta_0 + \beta_1 \text{AutoShock}_i + \beta_2 \text{AutoShockRhetoric}_i + \beta_3 \mathbf{Z}_i + \epsilon_i$. 95% confidence intervals are shown. Y_i represents: (a) aggregated preferences for the seven redistributive policies; and agreement with the statements: (b) “It would be unfair if the gap between those with high and low incomes were to increase as a result of workers being replaced by robots or artificial intelligence”; (c) “More jobs will be lost than gained as a result of automation by robots or artificial intelligence in the next 10 - 15 years”; (d) “It is likely that you or a member of your family will lose your job to robots or artificial intelligence in the next 10 - 15 years”. Binary variables are used to signal respondents’ exposure (1) or not (0) to each treatment. Controls are included for all variables listed in Table 2, except the education variables. Estimates should be interpreted as the change in percentage points relative to the base category. Significance of the marginal effect of the “Automation Shock with Rhetoric” treatment relative to the “Automation Shock” treatment is shown outside the square bracket on the right-hand side. All statistics presented in this figure are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning), disaggregated by respondents’ highest level of education. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6 Discussion and conclusions

As artificial intelligence and advanced robotics enable a growing number of tasks to be automated, some policy makers have called for a redistributive policy response designed to decouple growth in automation from growth in inequality. However, the findings of this study suggest that we should not assume that voters will necessarily support such redistributive policies - even as the negative labor market impacts of automation unfold.

This study has shown that as respondents perceive themselves to be more vulnerable to an automation shock, their support for redistributive policies tends to remain constant, and in some cases, *declines*. However, this study also shows that rhetoric that causes voters to perceive growth in automation-induced inequality as *unfair* is capable of preventing the decline in support for some redistributive policies and increasing support for others. The finding that respondents in possession of a Bachelor's degree are particularly responsive to the rhetorical treatment suggests that, as this influential group is increasingly exposed to automation, rhetoric could become an increasingly important factor in terms of the political viability of enacting a redistributive policy response to automation.

For researchers working in this rapidly expanding area of inquiry, this study underscores the value of looking beyond the labor market effects of automation as a determinant of political or policy preferences, to consider how the interaction with exposure to rhetoric influences voters' preferences. The study also highlights the value of conducting policy-level analysis. It reveals discrepancies between the way that preferences for specific redistributive policies respond to an automation shock, versus preferences for redistribution in more general terms. The policy-level analysis also reveals interesting nuance that suggests that vulnerable voters favor job-protecting policies over policies that would strengthen the social safety net in response to an automation shock.

Although some of the effects identified by this study are relatively small in absolute terms, it should come as no surprise that a single exposure to a brief passage warning of a threat that is "10 to 15 years" away has a small effect on policy preferences. Yet, we also know that first-hand experience of economic shocks - particularly labor market

shocks - exert large effects on individuals' political and policy preferences (Margalit 2019). The results presented throughout this paper therefore offer an indication of the direction that preferences may move in, how responses are likely to vary among different population sub-groups, and the mechanisms or conditions that could contribute to one response or another. While findings from Kerschbamer and Müller (2020) reassuringly suggest that distributional preferences predict voting behavior and political attitudes several years apart, further research to test the endurance of the effects identified in this study, over time and in the presence of competing rhetoric, would be beneficial.

References

- Aksoy, Cevat Giray, Berkay Ozcan, and Julia Philipp. 2020. “Robots and the gender wage gap”. *Working paper*.
- Alesina, Alberto F, Armando Miano, and Stefanie Stantcheva. 2020. *The polarization of reality*. Tech. rep. National Bureau of Economic Research.
- Alesina, Alberto, and Paola Giuliano. 2011. “Preferences for redistribution”. In *Handbook of Social Economics*, 1:93–131. Elsevier.
- Almås, Ingvild, Alexander W. Cappelen, and Bertil Tungodden. 2020. “Cutthroat Capitalism versus Cuddly Socialism: Are Americans More Meritocratic and Efficiency-Seeking than Scandinavians?” *Journal of Political Economy*. doi:10.1086/705551.
- Anelli, Massimo, Italo Colantone, and Pietro Stanig. 2018. “We Were the Robots: Automation in Manufacturing and Voting Behavior in Western Europe”. *Bocconi University*.
- Autor, David, Frank Levy, and Richard J Murnane. 2003. “The skill content of recent technological change: An empirical exploration”. *The Quarterly Journal of Economics* 118 (4): 1279–1333.
- Autor, David, and Anna Salomons. 2018. *Is automation labor-displacing? Productivity growth, employment, and the labor share*. Tech. rep. National Bureau of Economic Research.
- Baumberg, Ben. 2016. “The stigma of claiming benefits: a quantitative study”. *Journal of Social Policy* 45 (2): 181–199.
- Benjamini, Yoav, and Yosef Hochberg. 1995. “Controlling the false discovery rate: a practical and powerful approach to multiple testing”. *Journal of the Royal Statistical Society: Series B (Methodological)* 57 (1): 289–300.
- Berry, Mike. 2016. “The UK press and the deficit debate”. *Sociology* 50 (3): 542–559.
- Cavaille, Charlotte. *Implementing Blocked Randomization in Online Survey Experiments*. Available at <https://charlottecavaille.wordpress.com/blocked-randomization/>.

- Cavaillé, Charlotte, and Kris-Stella Trump. 2015. “The two facets of social policy preferences”. *The Journal of Politics* 77 (1): 146–160.
- Chong, Dennis, and James N Druckman. 2007. “Framing theory”. *Annual Review of Political Science* 10:103–126.
- Cochrane, Christopher, and Neil Nevitte. 2014. “Scapegoating: Unemployment, far-right parties and anti-immigrant sentiment”. *Comparative European Politics* 12 (1): 1–32.
- Dal Bó, Ernesto, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne. 2018. “Economic Losers and political winners: Sweden’s radical right”. *Unpublished manuscript, Department of Political Science, UC Berkeley*.
- DellaVigna, Stefano, and Matthew Gentzkow. 2010. “Persuasion: Empirical Evidence”. *Annual Review of Economics* 2:643–669.
- DellaVigna, Stefano, and Ethan Kaplan. 2007. “The Fox News effect: Media bias and voting”. *The Quarterly Journal of Economics* 122 (3): 1187–1234.
- Di Tella, Rafael, and Dani Rodrik. 2020. “Labour Market Shocks and the Demand for Trade Protection: Evidence from Online Surveys”. *The Economic Journal*. doi:10.1093/ej/ueaa006.
- Druckman, James N. 2001. “The implications of framing effects for citizen competence”. *Political Behavior* 23 (3): 225–256.
- ESS. 2016. *European Social Survey Round 8 Data. Data file edition 2.1. NSD - Norwegian Centre for Research Data, Norway – Data Archive and distributor of ESS data for ESS ERIC*. Available at 10.21338/NSD-ESS8-2016.
- Fernández-Albertos, Pepe, Aina Gallego, Alex Kuo, and Dulce Manzano. 2020. “Technological change, trade, and preferences for redistribution: Experimental evidence”. *Working paper*.
- Fisman, Raymond, Keith Gladstone, Ilyana Kuziemko, and Suresh Naidu. 2017. “Do Americans want to tax capital? Evidence from online surveys”. *National Bureau of Economic Research*.

- Frey, Carl Benedikt, Thor Berger, and Chinchih Chen. 2018. “Political machinery: did robots swing the 2016 US presidential election?” *Oxford Review of Economic Policy* 34 (3): 418–442.
- Frey, Carl Benedikt, and Michael A Osborne. 2014. “Agiletown: The relentless march of technology and London’s response”. *London Futures*.
- . 2017. “The future of employment: How susceptible are jobs to computerisation?” *Technological Forecasting and Social Change* 114:254–280.
- Frey, Carl Benedikt, Michael Osborne, Craig Holmes, E Rahbari, R Garlick, G Friedlander, G McDonald, E Curmi, J Chua, and P Chalif. 2016. “Technology at work v2.0: The future is not what it used to be”. *CityGroup and University of Oxford*.
- Georgiadis, Andreas, and Alan Manning. 2012. “Spend it like Beckham? Inequality and redistribution in the UK, 1983–2004”. *Public choice* 151 (3-4): 537–563.
- Gingrich, Jane. 2019. “Did State Responses to Automation Matter for Voters?” *Research & Politics* 6 (1). doi:10.1177/2053168019832745.
- Greszki, Robert, Marco Meyer, and Harald Schoen. 2015. “Exploring the effects of removing “too fast” responses and respondents from web surveys”. *Public Opinion Quarterly* 79 (2): 471–503.
- Im, Zhen Jie, Nonna Mayer, Bruno Palier, and Jan Rovny. 2019. “The “losers of automation”: A reservoir of votes for the radical right?” *Research & Politics* 6 (1). doi:2053168018822395.
- Ipsos. 2018. “AI, Automation and Corporate Reputation”. *Ipsos Knowledge Centre*.
- Karadja, Mounir, Johanna Mollerstrom, and David Seim. 2017. “Richer (and holier) than thou? The effect of relative income improvements on demand for redistribution”. *Review of Economics and Statistics* 99 (2): 201–212.
- Kerschbamer, Rudolf, and Daniel Müller. 2020. “Social preferences and political attitudes: An online experiment on a large heterogeneous sample”. *Journal of Public Economics* 182:104076.

- Korpi, Walter, and Joakim Palme. 1998. “The paradox of redistribution and strategies of equality: Welfare state institutions, inequality, and poverty in the Western countries”. *American sociological review*: 661–687.
- Kurer, Thomas, and Aina Gallego. 2019. “Distributional consequences of technological change: Worker-level evidence”. *Research & Politics* 6 (1). doi:2053168018822142.
- Kuziemko, Ilyana, Michael I Norton, Emmanuel Saez, and Stefanie Stantcheva. 2015. “How elastic are preferences for redistribution? Evidence from randomized survey experiments”. *American Economic Review* 105 (4): 1478–1508.
- Labour Party. 2019. “Alternative models of ownership”. <https://labour.org.uk/wp-content/uploads/2017/10/Alternative-Models-of-Ownership.pdf>.
- Margalit, Yotam. 2019. “Political responses to economic shocks”. *Annual Review of Political Science*.
- Meltzer, Allan H, and Scott F Richard. 1981. “A rational theory of the size of government”. *Journal of Political Economy* 89 (5): 914–927.
- Moore, Ryan T, and Keith Schnakenberg. 2016. “blockTools: Block, assign, and diagnose potential interference in randomized experiments”. *R package version 0.5-3*.
- NatCen Social Research. 2019. *British Social Attitudes Survey, 2017. [data collection]. UK Data Service*. Available at beta.ukdataservice.ac.uk/datacatalogue/series/series?id=200006.
- ONS. 2011. *2011 Census*. <https://www.ons.gov.uk/census/2011census/2011censusdata>.
- Owen, Erica. 2020. “Firms vs. Workers? The Politics of Openness in an Era of Global Production and Automation”. *Working Paper*.
- Palan, Stefan, and Christian Schitter. 2018. “Prolific. ac—A subject pool for online experiments”. *Journal of Behavioral and Experimental Finance* 17:22–27.
- Peer, Eyal, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. “Beyond the Turk: Alternative platforms for crowdsourcing behavioral research”. *Journal of Experimental Social Psychology* 70:153–163.

- Phillips, Daniel, John Curtice, Miranda Phillips, and Jane Perry. 2018. “British Social Attitudes: The 35th Report”. *The National Centre for Social Research*.
- Rehm, Philipp, Jacob S Hacker, and Mark Schlesinger. 2012. “Insecure alliances: Risk, inequality, and support for the welfare state”. *American Political Science Review* 106 (2): 386–406.
- Romer, Thomas. 1975. “Individual welfare, majority voting, and the properties of a linear income tax”. *Journal of Public Economics* 4 (2): 163–185.
- Shiller, Robert J. 2017. “Narrative economics”. *American Economic Review* 107 (4): 967–1004.
- Suiter, Jane. 2016. “Post-truth politics”. *Political Insight* 7 (3): 25–27.
- Thewissen, Stefan, and David Rueda. 2019. “Automation and the welfare state: Technological change as a determinant of redistribution preferences”. *Comparative Political Studies* 52 (2): 171–208.
- Van Oorschot, Wim. 2006. “Making the difference in social Europe: deservingness perceptions among citizens of European welfare states”. *Journal of European Social Policy* 16 (1): 23–42.
- Van Reenen, John. 2011. “Wage inequality, technology and trade: 21st century evidence”. *Labour Economics* 18 (6): 730–741.
- Webb, Michael. 2019. “The Impact of Artificial Intelligence on the Labor Market. Available at SSRN: <https://ssrn.com/abstract=3482150>”.
- Wu, Nicole. 2019. “Misattributed Blame? Attitudes Towards Globalization in the Age of Automation”. *Manuscript in preparation*.
- Yang, Andrew. 2020. “What is the Freedom Dividend?” <https://www.yang2020.com/what-is-freedom-dividend-faq/>.
- Zhang, Baobao. 2019. “No Rage Against the Machines: Threat of Automation Does Not Change Policy Preferences. MIT Political Science Department Research Paper. Available at SSRN: <https://ssrn.com/abstract=3455501>”.

A Appendix

Survey Instrument: Experimental treatments

Subjects were randomly assigned (with equal probability) to receive one of the following three informational treatments.

Automation Shock treatment

It is very important for the success of our research that you read the following passage carefully before continuing.

Experts predict that robots and artificial intelligence will destroy 1 in 3 jobs in the UK in the next 10 -15 years.

Many different types of workers are expected to be replaced by intelligent machines or computer programmes, including shop assistants, restaurant workers, warehouse staff, and administrators. For the first time, people with a university education are also thought to be at risk. Some accountants, lawyers, teachers and medical specialists have already begun to be replaced.

With so many jobs expected to be lost going forward many workers could be forced to accept lower wages or face long periods of unemployment.

We need to do something, but it's not clear what. None of the major political parties has set out a concrete plan of action to address the issue yet.

Do you personally know anyone employed in any of the jobs mentioned as being at risk?
Select ALL that apply: [The ordering of response options is randomized]

- Shop assistant
- Restaurant worker
- Warehouse worker
- Administrator
- Accountant
- Lawyer
- Teacher
- Medical specialist
- None of the above

Automation Shock with Rhetoric treatment

It is very important for the success of our research that you read the following passage carefully before continuing.

Experts predict that robots and artificial intelligence will destroy 1 in 3 jobs in the UK in the next 10 -15 years.

Many different types of workers are expected to be replaced by intelligent machines or computer programs, including shop assistants, restaurant workers, warehouse staff, and administrators. For the first time, people with a university education are also thought to be at risk. Some accountants, lawyers, teachers and medical specialists have already begun to be replaced.

With so many jobs expected to be lost going forward many workers could be forced to accept lower wages or face long periods of unemployment.

Some say that if corporations replace employees with machines, billionaire bosses and shareholders will profit while ordinary workers lose their jobs. But if we find a way to share the wealth that's created we could all work less and spend more time doing the things we enjoy.

Do you personally know anyone employed in any of the jobs mentioned as being at risk? Select ALL that apply: [The ordering of response options is randomized]

- Shop assistant
- Restaurant worker
- Warehouse worker
- Administrator
- Accountant
- Lawyer
- Teacher
- Medical specialist
- None of the above

Placebo

It is very important for the success of our research that you read the following passage carefully before continuing.

In recent decades, the types of jobs that people in the UK have been employed to do has changed a lot.

Some of the occupations that used to be very common are much less commonplace today. For instance, milkman, weaver, and TV or radio repairer were all typical occupations in the past, but are among the least common jobs in the UK today.

The most common occupations today are jobs in the service sector, such as shop assistants, administrators, nurses and care workers. Although these occupations have been popular for many decades, today more people are employed in the service sector than ever before.

Looking forward to the next 10 – 15 years, experts predict that this trend will continue with an even larger segment of the working population expected to be employed in the service sector in the years to come.

Do you personally know anyone employed in any of the jobs mentioned in the passage above? Select ALL that apply: [The ordering of response options is randomised]

- Milkman/Milkwoman
- Weaver
- TV repairer
- Radio repairer
- Shop assistant
- Administrator
- Nurse
- Care worker
- None of the above

Additional control treatment

To test for the possibility that the original Placebo treatment influenced preferences for increasing income support or increasing the minimum wage, the survey was re-administered to a sample of 300 respondents, using the following control text.

Control text

It is very important for the success of our research that you read the following passage carefully before continuing.

For millennia, our ancestors have been collecting and consuming honey. The healing properties of honey are mentioned in the folklore of the Ancient Greeks and Egyptians, and honey is counted as an important ingredient in traditional medicine.

Honey was traditionally thought to cure ailments such as gastric disturbances, ulcers, skin wounds, and skin burns. Today scientists have found evidence suggesting there may be some truth to the “ancient wisdom”.

Some studies suggest that a layer of honey on a wound or burn can aid recovery more quickly than antiseptic and gauze, though evidence of this is mixed. Others suggest that consuming honey mixed with warm water and lemon can be an effective alternative to over the counter cough remedies.

Policy preference questions

Thinking about the expected changes to the UK job market that you have just read about, would you agree or disagree with the government implementing the following policies going forward?

Going forward, the government should...

- a) ...increase benefits paid to people with low incomes.
 - Strongly agree
 - Agree
 - Somewhat agree
 - Neither agree nor disagree
 - Somewhat disagree
 - Disagree
 - Strongly disagree
- b) ...increase the minimum wage that employers can legally pay their staff.
- c) ...increase taxes on corporations' profits.
- d) ...increase taxes on incomes of over £150,000.
- e) ...increase benefits paid to people who are unemployed.
- f) ...pay every citizen enough to cover the cost of basic necessities, regardless of whether they work or not.
- g) ...make it possible for workers to reduce their working hours by 20% with no reduction in pay.
- h) ...increase restrictions on the number of immigrants allowed to work in the UK.

Expectations and beliefs questions

In this final section, you will be asked about your views and expectations. If you are unsure of how to answer at any time, please give your best guess.

Do you agree or disagree with the following statements?

a) More jobs will be lost than gained as a result of automation by robots or artificial intelligence in the next 10 - 15 years.

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

b) It is likely that you or a member of your family will lose your job to robots or artificial intelligence in the next 10 – 15 years.

c) It would be unfair if the gap between those with high and low incomes were to increase as a result of workers being replaced by robots or artificial intelligence.

d) In general, the major political parties in the UK can be trusted to deliver on the promises they make.

e) Benefits for unemployed people are too high and discourage people from working.

f) Increasing taxes on corporations usually leads to job losses.

Table 3: Effect of treatments on policy preferences

	Increase Income Support	Increase Minimum Wage	Increase Corporation Tax	Increase Top Income Tax	Increase Unemployment Benefit	Provide a UBI	Reduce Working Time	Restrict Immigration
<i>Marginal effects of treatments relative to baseline</i>								
Automation Shock	-3.81*** (1.40)	-2.79** (1.21)	0.47 (1.34)	2.17 (1.51)	2.50 (1.54)	3.70** (1.69)	6.04*** (1.64)	1.98 (1.65)
Automation Shock with Rhetoric	-0.49 (1.35)	0.68 (1.15)	4.33*** (1.24)	4.45*** (1.48)	3.65** (1.52)	3.58** (1.70)	10.24*** (1.61)	0.96 (1.66)
<i>Marginal effects between treatments</i>								
Auto Shock with Rhetoric relative to Auto Shock	3.32** (1.42)	3.48*** (1.22)	3.87*** (1.32)	2.28 (1.48)	1.15 (1.56)	-0.12 (1.73)	4.20*** (1.60)	-1.02 (1.69)
<i>Controls</i>								
Male	-0.38 (1.15)	-3.79*** (0.98)	0.41 (1.08)	1.18 (1.23)	5.66*** (1.27)	1.96 (1.40)	-0.36 (1.34)	-1.39 (1.37)
Age 35 - 54	1.61 (1.51)	0.05 (1.29)	4.14*** (1.39)	4.91*** (1.66)	4.50** (1.78)	-0.35 (1.81)	1.58 (1.82)	4.59** (1.86)
Age 55+	1.64 (1.71)	-2.85* (1.46)	1.70 (1.59)	4.69** (1.85)	4.41** (1.96)	-0.86 (2.04)	1.42 (1.93)	8.64*** (2.02)
Asian	5.52*** (2.03)	0.96 (2.14)	-1.93 (2.19)	-1.10 (2.35)	1.69 (2.70)	1.73 (2.60)	3.06 (3.00)	-1.45 (2.79)
Black	5.61* (2.92)	6.16*** (2.10)	-2.08 (2.98)	-8.01** (3.44)	2.82 (3.26)	-4.43 (3.71)	3.07 (3.87)	-8.87** (4.24)
Other non-white	4.23 (3.68)	-0.31 (3.27)	-1.82 (3.01)	-4.44 (3.82)	4.57 (3.73)	-0.91 (4.09)	-5.39* (3.09)	-16.03*** (3.75)
Medium income (£20-60k)	-6.90*** (1.38)	-3.93*** (1.18)	-0.89 (1.34)	-1.58 (1.49)	-8.70*** (1.55)	-7.68*** (1.68)	-5.48*** (1.61)	-1.75 (1.74)
High income (£60k+)	-13.35*** (1.92)	-6.55*** (1.62)	-3.59** (1.83)	-8.38*** (2.15)	-15.44*** (2.08)	-11.36*** (2.30)	-6.41*** (2.20)	-5.09** (2.24)
Unemployed	0.97 (2.57)	0.52 (2.16)	-2.58 (2.66)	-7.33** (2.92)	10.01*** (3.09)	4.13 (2.89)	-0.49 (2.96)	-0.76 (3.27)
Student	2.65 (2.74)	-3.97 (2.72)	-0.70 (3.16)	-1.43 (3.44)	8.24*** (3.13)	7.63** (3.09)	-3.31 (3.12)	-15.65*** (3.59)
Not in labour force	0.86 (1.47)	-0.44 (1.21)	1.27 (1.31)	-0.88 (1.52)	5.28*** (1.57)	1.12 (1.81)	-7.47*** (1.67)	-2.71 (1.72)
Has a degree	-1.15 (1.04)	-1.07 (0.90)	1.21 (0.96)	1.84* (1.11)	3.78*** (1.17)	1.23 (1.28)	-0.76 (1.24)	-13.23*** (1.29)
Conservative	-10.24*** (1.68)	-7.87*** (1.53)	-8.70*** (1.63)	-11.52*** (1.84)	-13.22*** (1.83)	-17.15*** (2.09)	-6.92*** (1.97)	16.58*** (1.94)
Labour	8.88*** (1.50)	6.73*** (1.27)	6.78*** (1.33)	6.81*** (1.56)	8.51*** (1.72)	7.22*** (1.79)	8.23*** (1.76)	-1.75 (1.96)
Placebo mean	67	79	77	73	50	60	51	55
Observations	2,296	2,296	2,296	2,296	2,296	2,296	2,296	2,296

Notes: Estimates from: $\text{PolicyPref}_i = \beta_0 + \beta_1 \text{AutoShock}_i + \beta_2 \text{AutoShockRhetoric}_i + \beta_3 \mathbf{Z}_i + \epsilon_i$. Robust standard errors are presented in parentheses. The outcome variables are coded on a seven-point scale standardized from zero (strongly disagree) to 100 (strongly agree) and signal agreement with the statements: “Going forward, the government should... (i) increase benefits paid to people with low incomes; (ii) increase benefits paid to people who are unemployed; (iii) increase the minimum wage that employers can legally pay their staff; (iv) increase taxes on incomes of over £150,000; (v) increase taxes on corporations’ profits; (vi) make it possible for workers to reduce their working hours by 20% with no reduction in pay; (vii) pay every citizen enough to cover the cost of basic necessities, regardless of whether they work or not; (viii) increase restrictions on the number of immigrants allowed to work in the UK. Binary variables are used to signal respondents’ exposure (1) or not (0) to each treatment. Controls are binary. The following are left out as reference categories: Non-male; Age 18 - 34; White; Low income (<£20k); Employed; No degree; Other party. Estimates should be interpreted as the change in percentage points, relative to the baseline. All statistics presented in this table are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning). ***p<0.01, **p<0.05, *p<0.1.

Table 4: Summary statistics and policy preferences from the original Placebo group and the follow-up Control group

	Original Placebo	Follow-up Control
<i>Policy preferences (mean)</i>		
Increase Income Support	66.8	65.8
Increase Minimum Wage	78.8	78.8
<i>Socio-demographics (%)</i>		
Male	48.8	48.7
Age 18 - 34	28.5	28.6
Age 35 - 54	35.7	38.2
Age 55+	35.8	33.2
White	87.1	87.1
Asian	6.4	5.8
Black	3.3	2.1
Other ethnicity	3.2	5.1
Employed (full or part)	61.1	63.7
Unemployed	5.0	4.7
Student	5.3	8.7**
Has university degree	27.9	27.9
Conservative voter	28.9	28.9
Labour voter	32.1	32.1
Observations	764	300

Notes: Column 1 displays weighted summary statistics from the Placebo group from the original survey experiment. Column 2 displays weighted summary statistics from a follow up survey experiment using an alternative control text. Asterisks in column 2 indicate significant p-values from t-tests comparing the weighted means of the two samples. ***p<0.01, **p<0.05 , *p<0.1.

Figure 8: Effect of treatments on policy preferences by education-level

Notes: Estimates from: $\text{PolicyPref}_{i,e} = \beta_0 + \beta_1 \text{AutoShock}_{i,e} + \beta_2 \text{AutoShockRhetoric}_{i,e} + \beta_3 \mathbf{Z}_{i,e} + \epsilon_{i,e}$. 95% confidence intervals are shown. The outcome variables are coded on a seven-point scale standardized from zero (strongly disagree) to 100 (strongly agree) and signal agreement with the statements: “Going forward, the government should... (i) increase benefits paid to people with low incomes; (ii) increase the minimum wage that employers can legally pay their staff; (iii) increase taxes on incomes of over £150,000; (iv) increase taxes on corporations’ profits; (v) increase benefits paid to people who are unemployed; (vi) pay every citizen enough to cover the cost of basic necessities, regardless of whether they work or not; (vii) make it possible for workers to reduce their working hours by 20% with no reduction in pay”; (viii) increase restrictions on the number of immigrants allowed to work in the UK”. Binary variables are used to signal respondents’ exposure (1) or not (0) to each treatment. All variables listed in Table 2 are included as controls, except the education variable. All statistics presented in this figure are estimated using the weighted regression sample (the 92% of the sample remaining after cleaning), disaggregated by respondents’ highest completed level of education.