

# Attitudes toward automation and the demand for policies addressing job loss: the effects of information about trade-offs

Beatrice Magistro<sup>\*,‡</sup>, Peter John Loewen<sup>\*</sup>, Bart Bonikowski<sup>†</sup>, Sophie Borwein<sup>§</sup>, and Blake Lee-Whiting<sup>\*</sup>

<sup>\*</sup> University of Toronto

<sup>†</sup> New York University

<sup>§</sup> Simon Fraser University

<sup>‡</sup> *corresponding author*, [beatrice.magistro@utoronto.ca](mailto:beatrice.magistro@utoronto.ca)

## Abstract

Does providing information about the costs and benefits of automation affect support for automation and for different policies in response? To answer this question, we use a combination of survey and conjoint experiments across four advanced economies (Australia, Canada, the UK, and the US). Our results show that despite people’s relatively fixed policy preferences, their evaluation of automation—and therefore potentially the issue’s political salience—is sensitive to information about its trade-offs, especially to price changes. This suggests that automation may have different political consequences depending on how it is framed by the media and political actors.

## 1 Introduction

Modern economies are experiencing major changes, including labour dislocation, as a result of a new wave of automation and artificial intelligence (AI). While automation is likely to lead to aggregate efficiency gains, it also has distributional consequences, as some workers will lose their jobs while others will get new jobs (Aghion et al., 2020, 2021; Autor & Dorn, 2013; Autor et al., 2003; Bessen, 2019). As a result of the heterogeneity of these labour market disruptions, there is also a growing literature on the political consequences of

automation and AI (Anelli et al., 2019; Di Tella & Rodrik, 2020; Frey & Osborne, 2017; Gallego et al., 2022; Im et al., 2019). Much of the literature finds that workplace automation matters for voting behavior (Anelli et al., 2019; Im et al., 2019; Kurer, 2020), with losers of technological change more likely to vote against the political establishment (whether they lean left or right depends on how they understand technological threat (Borwein et al., 2022)). However, the literature provides more mixed findings on whether exposure to automation-related risk influences the policy preferences of workers: automation risk is not always related to preferences for more redistribution or for policies protecting jobs (Gallego & Kurer, 2022).

Much of the existing literature draws upon standard political economy models, according to which labour market risks affect economic interest, which in turn affects political preferences and voting behavior. But are people aware of the trade-offs of automation for labor markets, economic growth, and the prices of goods and services? If so, how do they evaluate those trade-offs? This is particularly important because economic shocks have different political consequences depending on how they are perceived by voters and interpreted by political actors.

There are reasons to believe that voters may not be aware of the potential effects of automation. First, although technological change has been the largest determinant of job polarization<sup>1</sup>, modern political parties have not consistently or widely mobilized voters against technological change; instead they have turned to simpler explanations for recent structural changes in labour markets, like trade and immigration, that tend to fuel out-group resentments (Gallego & Kurer, 2022; Mutz, 2021). At the same time, despite the evident economic benefits stemming from technological change, automation is generally discussed in a negative light by the media as well as by the scholarly literature, as it is more often framed as being associated with job loss rather than job creation (Anelli et al., 2019; Frey & Osborne, 2017; Im et al., 2019). Existing

---

<sup>1</sup>By job polarization we mean the relative growth of employment in high-skill jobs and low-skill jobs amid the concurrent decline in middle-skill jobs (Autor & Dorn, 2013; Goos et al., 2014)

experimental research on automation similarly relies on treatments that only communicate the costs of technological change, for instance, stressing the number of jobs that may be lost (Di Tella & Rodrik, 2020; Jeffrey, 2021; Jeffrey & Matakos, 2021; Mutz, 2021; Werfel et al., 2022), or communicate trade-offs only in vague terms, without providing any hypothetical numbers (Gallego et al., 2021).

In this paper, we examine opinions towards automation more broadly. In particular, we ask: Does providing information about the trade-offs of automation affect support for automation and for different policies in response? In our treatments, we provide information not only about potential job and wage losses, but also about potential job and wage gains, and more importantly, potential price changes as a result of new technology.

We use a combination of survey and conjoint experiments across four advanced, liberal market economies (Australia, Canada, the UK, and the US), with a total sample of about 8,000 individuals, to answer these questions. Each respondent is randomized into either a costs-only news article treatment, a low information trade-off treatment, or a high information trade-off conjoint treatment. In the news article treatment individuals read, in sensational terms, about only the costs of automation, i.e., the job and wage losses, but are not given information about the benefits. Individuals in the low information group read about a firm introducing a new computer-based productivity improving technology and that this will lead to some jobs gains, some job losses, and possibly lower prices, but are not given any exact numbers. In the high information conjoint treatment each individual sees a combination of four random tables, each with varying quantified costs and benefits of automation (in terms of final products' price changes and changes in the number of employed workers and their wages). Finally, after the respective vignettes, individuals in all groups are asked how much they support the decision to automate, and what kind of policies in response to automation they favor governments taking.

Our analysis proceeds in three steps. First, we analyze results from a between-subjects experiment with three groups: the news treatment, the low

information treatment, and the high information conjoint (for the latter we take the average of each respondent's four answers to the four tables they see on trade-offs). We hence examine respondents' attitudes toward automation (whether they think it is fair for the firm to automate, and whether they would make the same decision if they were the CEO of the company in question), and toward policy responses to it (social spending, basic income, job guarantee, unskilled immigration restrictions, skilled immigration restrictions, trade restrictions, retraining workers, and taxing automation).

Second, we use conjoint analysis to analyze respondents' attitudes toward different automation scenarios. This tells us what attributes (type of product, number of high skilled workers, number of low skilled workers, wages of high skilled workers, wages of low skilled workers, price change of product) causally increase or decrease support for automation on average when varied independently of the other attributes included in the design. Furthermore, within the conjoint experiment, we measure whether individuals understand the numbers they see. We hence build a measure of knowledge and examine whether people who correctly computed the costs and benefits of the new innovation have different attitudes towards automation.

We find that relative to the costs-only news information treatment, people in the low information group, who read about the trade-offs of automation in vague terms, display significantly more positive views of automation. Respondents in the high information group, who see the exact numbers of winners and losers across different scenarios, on average also show more positive views of automation than people in the costs-only news condition<sup>2</sup>, but this increase is not as large as for people in the low information group. This suggests that people in the low information treatment group, absent specific details about winners and losers, may be envisioning a scenario in which the net gains of automation are larger relative to the average conjoint scenario.

When looking at support for different policies in response to automation,

---

<sup>2</sup>This is true of fairness, but not of whether they would do the same thing if they were CEO. For the latter the effects of the high information conjoint treatment and the news treatment are not significantly different from each other.

however, there are no discernible differences across the three treatment groups, suggesting that policy preferences are relatively sticky. Overall, individuals in all three groups favor retraining individuals affected by automation more than any other policy option. The least popular policy responses are restricting skilled immigration and trade.

Finally, our conjoint experiments presented subjects with information on job numbers and wage losses or gains across high and low-skilled employees, as well as price changes. Our results show that respondents are most sensitive to price changes. When prices decrease by 20% to 50% as a result of the new technology, respondents are much more likely to think automation is fair and that they would make the same decision if they were the CEO than when there is no price change. In a similar way, when fewer (more) workers are lost (gained), they are more supportive of automation than when more (fewer) workers are lost (gained). They are also sensitive to varying declines in the wages of workers losing out from automation, but not to increases in the wages of workers winning from automation. Respondents do not appear to be sensitive to the type of firm or product in question, whether an auto company producing cars, an aviation firm producing planes, an electronics company making smartphones, or a pharmaceutical company producing vaccines.

Our findings have several implications. First, relative to a fear-mongering media message that only emphasizes the costs of automation, information on its trade-offs can increase support for automation significantly. This is especially true of generic information; more specific information leads to lower increases in support for automation than vague information on average, but these estimates vary considerably based on the relative gains and losses from automation, as highlighted by our conjoint analysis. Furthermore, people's policy preferences are relatively stable—none of the information treatments shift respondents' preferences over policy responses, as they are consistently more supportive of retraining workers relative to other policy options. Finally, no previous experimental study factors in potential price changes among the potential benefits or costs of automation. We find that this is the feature to which respondents are

by far most sensitive. While respondents also react to changes in the number of workers and less so to changes in wages, when prices decrease by 20% or 50% they are much more likely to support automation. In the latter scenario, the point estimate is approximately indistinguishable from that of the low information group. These findings suggest that the political consequences of automation may differ based on how and if political actors decide to politicize the issue and on the information they provide to voters.

## 2 The Political Economy of Automation

Labour markets may be profoundly reshaped by technological change, in particular by automation and artificial intelligence. While many argue that these developments will produce mass unemployment as new technologies replace labour, others stress how both historical and more recent evidence suggest that these fears are unwarranted. The big technological revolutions that have produced the largest increases in growth and prosperity—the steam engine in the early 1800s and electricity in the 1920s—did not produce the mass unemployment that was anticipated by some (Aghion et al., 2021). The literature has evolved from a more negative view of automation, seen as primarily destroying jobs, to a more positive view of automation as productivity enhancing (Aghion et al., 2020, 2021; Bessen, 2019). This latter approach emphasizes the direct productivity effect of automation: firms that automate become more productive, which allows them to lower their quality-adjusted prices, and in turn to increase demand for their products. The resulting effect is an increase in employment by these firms.

Separate from the debate about the aggregate effects of automation on employment, there is also a debate on its distributional consequences. In particular, recent works show that job losses have been concentrated in occupations that feature routine tasks, as opposed to those requiring human interaction and higher education, which are least at risk (Autor & Dorn, 2013; Autor et al., 2003). Some studies find evidence of a reallocation of workers between oc-

cupations (Humlum, 2019), with labour demand shifting from low-skilled to high-skilled workers. Kurer and Gallego (2019) focus on the aggregate decline in routine work as a result of technological change and find that many routine workers manage to keep their jobs until early retirement, and that the decline is mostly driven by higher exit rates and lower entry rates rather than layoffs. Although the job risk estimates vary, experts agree that automation and AI will continue to transform the nature of work (Autor et al., 2020). Some workers will lose their jobs to automation, others will get new jobs, and many will need to acquire new skills to transition across occupations.

As a result of the heterogeneity of these labour market disruptions, there is also a growing literature on the political consequences of technological change. A still relatively small literature investigates the effects of technological change on vote choice, and an even more limited literature looks at citizens' preferred policies to address technological change. Recent findings suggest that higher subjective and/or objective risk of job loss from automation is related to both left and populist right voting (Anelli et al., 2019; Borwein et al., 2022; Frey et al., 2018; Im et al., 2019). This literature mostly focuses on left-behind voters, while neglecting the large majority of workers who benefit from innovation. One exception is Gallego et al. (2022), whose study of the active labour force in the UK between 1997 and 2017 provides evidence that technological adoption was economically beneficial for workers with middle and high levels of education but produced small negative effects for low education workers. Furthermore, growth in automation increased support for the incumbent party and voter turnout among those who benefited from technological change.

The growing literature on automation and policy preferences has been influenced by the existing scholarship on the labour market effects of other structural changes, such as trade or immigration. Much of this literature draws upon political economy theories. The expectation is that workers more negatively affected by technological change should be more likely to oppose automation and to favor compensatory redistributive or protectionist measures. However, while it has been established that technological change is the most important

structural driver of job polarization—and possibly income inequality—in recent years, economic losers do not appear to directly or explicitly blame technological change for their relative decline in economic well-being (Gallego & Kurer, 2022). Specifically, the literature on automation and political preferences has investigated if automation risk impacts preferences for redistribution and other economic attitudes, including support for more spending on active labour market policies, and the evidence is so far mixed (Busemeyer et al., 2022; Busemeyer & Sahn, 2021; Dermont & Weisstanner, 2020; Guarascio & Sacchi, 2021; Im, 2021; Jeffrey, 2021; Sacchi et al., 2020; Thewissen & Rueda, 2019). These mixed findings are possibly due to the fact that different studies use different measures and model specifications, but may also vary due to context or different political discourse. Alternatively, it is possible that workers at greater risk of automation are more likely to misattribute their relative decline in economic well-being to trade and immigration (Wu, 2021), which in turn, may lead to their increased support for reactionary policy responses, such as restricting trade or immigration, rather than less distortionary policies, such as retraining workers or providing unemployment insurance for displaced workers (Jaimovich et al., 2020). In particular, policies directly related to technological change are rarely discussed by political parties (Rodrik, 2018). While populist parties campaign against trade and immigration, we do not see the same trend against technology or automation. Hence, it is plausible that discontent caused by technological change may manifest itself in the political arena through debate on other issues, such as trade and immigration, which offer clear out-groups to mobilize against and more straightforward policies to counteract them.

Overall, individuals may not be able to distinguish the role of technological change from other structural changes in affecting their economic well-being or understand the trade-offs involved with automation, and as a result they may be uncertain about how they feel about automation or which policies are more likely to favor in response to it. Our main contribution is to answer a series of key questions: Does providing information about the trade-offs of automation affect support for automation and for different policies in response to it? How



do people view automation in the absence of detailed information on its effects? When information is indeed provided, how do people weigh the different trade-offs? Under what conditions are they more or less likely to favor automation and to support retraining workers or providing unemployment insurance for displaced workers as opposed to protectionism?

### 3 Theoretical Expectations

The existing literature on automation and similar structural labour market changes enables us to make some broad predictions about the results of our experiments.<sup>3</sup> We summarize these below before turning to our data and methods.

First, we know that automation tends to be portrayed in a negative light by the media as well as by the scholarly literature, as it is more often associated with job loss than job creation (Anelli et al., 2019; Frey & Osborne, 2017; Im et al., 2019). A case in point is the prediction by Frey and Osborne (2017) that 47% of current jobs could be lost to automation in the next 20 years, which has attracted significant scholarly and media interest (Gallego et al., 2021). Based on this, we expect that a news article treatment highlighting only the costs of automation would be consistent with the kind of information people are routinely exposed to in the media. Relative to such costs-only accounts, we expect that treatments consisting of generic or specific information about the trade-offs of automation—highlighting both its costs and benefits—would lead to more positive views on automation. Specifically, we expect the increase in support for automation to be larger among people who receive more specific, as opposed to more generic, information about the trade-offs of technological change. The rationale behind this expectation is that people assigned to the low information trade-off condition, who are not shown any numerical estimates of costs and benefits, will likely overestimate the relative costs of automation.

---

<sup>3</sup>We registered our pre-analysis plan on OSF (Magistro et al., 2022). We have diverted somewhat from our PAP where we felt that doing so would make the presentation of the data clearer—we changed the framing but not the predictions.

This would result in a smaller increase in support for automation compared to the increase predicted for the high information group. We expect that a high information treatment group that sees the specific costs and benefits of innovation would get a more accurate picture and hence have more positive views of automation, as they update their priors.

The literature about labour market risks suggests that in response to higher risk of unemployment or income loss, people are more likely to support redistributive policies (Iversen & Soskice, 2001; Thewissen & Rueda, 2019). However, Colantone and Stanig (2018) note that when it comes to trade, protectionist policies are often favored over compensatory ones, as they allow workers to preserve their jobs, although at a much larger cost to society. When it comes to technology, workers may favor policies that attempt to prevent or disincentivize the adoption of new technologies (Abbott & Bogenschneider, 2018; Dauth et al., 2021). Which of these policies in response to automation are people more likely to support? We expect that this may depend on the information they see, i.e., which combination of costs and benefits the new technology would give rise to. In response to job losses due to automation, people exposed to quantified costs and benefits should be more likely to support policies involving active labour market interventions (such as re-training) and redistributive social spending (such as a basic income or unemployment insurance), rather than policies that protect jobs or inhibit automation (such as halting automation, taxing it, or restricting trade or immigration) compared to individuals in the more generic low information group and to those exposed to the news article treatment (the latter group should be the most supportive of protectionist policies). The rationale is that in scenarios in which the costs of automation are more salient than its benefits, people may display more protectionist attitudes in an effort to directly save jobs rather than compensatory or redistributive views that offset job loss instead of preventing it.

Second, we also analyze respondents' attitudes towards automation using conjoint analysis. We investigate which attributes, on average, causally increase or decrease support for automation when varied independently of the other at-

tributes included in the design. The attributes we vary include type of product manufactured by the company, number of high skilled jobs gained, number of low skilled jobs lost, changes in the wages of high skilled workers, changes in the wages of low skilled workers, and the estimated price change of products manufactured by the focal company. We are agnostic as to which of the attributes may matter most, but we expect respondents to respond rationally to changes within attributes: for example, as prices decrease as a result of technological innovation, respondents should become more supportive of automation, all else equal. In line with prospect theory (Kahnemann & Tversky, 1979), it is also possible that job and wage losses may be weighted more heavily than job and wage gains. These effects are likely to vary by respondents' economic literacy, which we measure in the high information conjoint treatment group by asking respondents to compute the costs and benefits of the new innovation. Magistro (2022) finds that people who correctly compute the costs and benefits of a policy are more likely to support or oppose the policy in question. Similarly, we expect that people who display higher economic literacy should be more responsive to changes in the attributes, since they should have a better understanding of the trade-offs involved.

## 4 Research design

What happens after a firm introduces new computer-based productivity improving technology? Theoretically, we can think of a few different scenarios, which we use to motivate our experimental treatments.

First, assuming we initially have a perfectly competitive market for the main good being produced by the firm in question, a new computer-based technology will make the firm more productive, shifting the supply curve out, which in turn will drive prices down for consumers. Provided that demand is elastic enough to prices, then product demand will increase, which will result in net job growth.<sup>4</sup> However, the expectation is that low-skilled and high-skilled

---

<sup>4</sup>We are describing process (i.e., reducing marginal costs) rather than product (i.e., in-

workers in the firm will be affected heterogeneously by a new computer-based productivity-improving technology. In this scenario, where demand is elastic enough to prices, after automation, a majority of high-skilled workers will benefit. These include workers performing certain technical skills, required to deploy, operate and maintain the new digital technologies, such as AI, big data, and machine learning specialists. However, a minority of low-skilled workers will lose their jobs, as automation will reduce the demand for jobs with more repetitive tasks that can be easily automated, such as assembly and factory workers (Centre for the New Economy and Society, 2018). In this hypothetical scenario, the gains for consumers, producers, and for high-skilled workers would be substantial enough that it would be more efficient to compensate the low-skilled workers through retraining programs, unemployment insurance, or other policy options, than to stop innovation altogether.

What if demand is inelastic? In this scenario, as prices go down as a result of new innovative technologies, demand does not increase enough to result in net job growth. Again, we could consider heterogeneous effects, where only a minority of high-skilled workers benefit, while a majority of low-skilled workers lose their jobs.

If we relax the assumption that the market is perfectly competitive, we could also envision a scenario in which the firm chooses not to decrease prices as a result of the innovation. Alternatively, even in a perfectly competitive market, there could be a scenario where prices are sticky in the short-run and hence they may not decrease right away.

Analyzing these different scenarios, we can envision different outcomes for high-skilled and low-skilled workers. We use a simplified version of the examples outlined here to build our survey experiment, varying the information

---

producing a new product) innovation. The former shifts the supply curve out and leads to lower prices, gains in consumer welfare, and depending on the elasticity of demand, it can lead to jobs gains or losses. The latter could possibly lead to even larger gains for consumers if the demand curve is shifted out drastically enough, as quality-adjusted prices may decrease even more than production costs and consumer surplus may be significantly larger. Product innovation is historically responsible for larger increases in consumer welfare than process innovation (Menaldo, Forthcoming), hence our example and related treatments likely provide a lower bound on the benefits of automation and innovation.

treatments to see how people react.

## 4.1 Survey experiment

To assess the role of cost-benefit information on attitudes toward automation and policies in response to it, we fielded an online survey experiment in Australia, Canada, the UK, and the US using the survey-sample provider Cint.<sup>5</sup> We chose four relatively similar countries, which share the same language, are essentially liberal market economies, have similar historical origins, and form the group of countries in a more advanced stage of de-routinization<sup>6</sup> that have adopted technology more intensively (De La Rica & Gortazar, 2016). We recruited respondents over 18 years old (we undersampled older individuals over 65 since we wanted a sample representative of working age individuals) and we applied quotas for age, gender, and region. Excluding respondents with questionable IP addresses, duplicate IDs, and very fast completion times (below the second percentile), a total of 8,033 respondents were surveyed: 1,955 in Australia, 1,972 in Canada, 2,031 in the UK and 1,966 in the US, with a median completion time of 15 minutes. In order to be able to conduct conjoint analysis we assigned 2/3 of respondents to the high information conjoint treatment group, 1/6 to the news treatment group, and 1/6 to the low information group. We estimate all models using ordinary least squares and pooling all four countries<sup>7</sup>. We also conduct analyses examining countries separately, but since with a few exceptions, respondents in different countries do not display different attitudes towards automation or demand different policies in response to job loss when presented with the same type of information, we only report country-by-country analyses in the Appendix. The following sections describe each treatment condition in detail.

---

<sup>5</sup>Fielding occurred between March 11 and 28, 2022. For more information on the survey-sample provider see <https://www.cint.com/>

<sup>6</sup>By de-routinization we mean the displacement of labour from routine tasks and hence greater reallocation of low-skill workers from routine task-intensive occupations to service occupations.

<sup>7</sup>We also run the same models controlling for an easy attention check and the results do not change. In the attention check question we simply ask respondents to select “Somehow disagree” among the response categories.

### **News article treatment group**

In the news article treatment vignette, individuals read, in sensational terms, about the costs of automation—with reference to job and wage losses—but were not given information about the benefits:

Assembly and factory jobs are at risk at a manufacturing firm in [Australia, Canada, the UK, the US], as management has decided to introduce a new computer-based productivity improving technology, which would lower production costs significantly. We interviewed an employee there for 20 years, who said that the technology shock will be devastating: “Up to 150 people will become unemployed and the rest would have to accept lower wages” he added.

After viewing the vignette, respondents were asked how much they agreed or disagreed with the following statements (from Borwein et al. (2021)), using a five-point scale from strongly agree to strongly disagree (with a sixth “don’t know” option):

- The company’s decision to introduce the new technology is fair.
- If you were the CEO of the company you would make the same decision to introduce the new technology.

Subsequently, they were asked how much they agreed or disagreed with each of the following government policies (from Borwein et al. (2021)), again on a five-point scale with a “don’t know” option:

- Expand social spending to support laid-off workers, and workers in similar positions.
- Implement a basic income that gives every adult a set amount of money from government on a regular basis.
- Pay to retain displaced workers and guarantee them jobs.

- Reduce the number of unskilled immigrants entering the country for work.
- Reduce the number of skilled immigrants entering the country for work.
- Restrict international competition by increasing trade barriers on goods and services to [Australia, Canada, the UK, the US].
- Fund programs to re-skill workers for new jobs.
- Directly tax companies that replace workers with machines and robots.

### **Low information group**

Individuals in the low information vignette were exposed to information about a new innovation which involves trade-offs, but were not given any precise estimates of the innovation's costs and benefits:

“A manufacturing firm in [Australia, Canada, the UK, the US] decides to introduce a new computer-based productivity improving technology. As a result of this innovation, production costs will decrease, and the price of the company's final products could also decrease. Furthermore, while some jobs will be gained, others will be lost.”

After viewing the vignette, respondents were asked the same questions as in the news article condition about their attitudes toward the automation decision and toward different policy responses to it.

### **High Information Conjoint Treatment Group**

Individuals in the conjoint group first saw a pre-treatment prompt:

A manufacturing firm in [Australia, Canada, the UK, the US] decides to introduce a new computer-based productivity improving technology. As a result of this innovation, production costs will decrease, and the price of the company's final products could also decrease.

This innovation could create new high-skilled jobs. These highly demanded high-skilled workers include those performing certain technical skills, required to deploy, operate and maintain the new digital technologies, specifically, AI, big data, and machine learning specialists. However, some low-skilled workers, specifically assembly and factory workers, who perform jobs with more repetitive tasks that can be easily automated, will lose their jobs to new machines and technologies. Furthermore, the remaining low-skilled workers will also see a cut in their yearly pay. The following tables show different possible scenarios.

Then, each individual saw a sequence of tables (see table 1 for an example), each with varying estimates of the costs and benefits of automation<sup>89</sup>.

After the first table respondents were asked the following multiple-choice questions to determine whether they understood the effects of the innovation:

- How much cheaper does the main product become after the innovation?  
[correct options are either 0%, 20% or 50%.]
- What is the total number of workers before the innovation? And after?

After viewing each table, respondents were asked how much they agreed or

---

<sup>8</sup>For simplicity, we keep one job per type (one high-skilled—AI, big data, and machine learning specialists—and one low-skilled: factory and assembly workers). We do not include a scenario where low-skilled workers actually benefit. Theoretically, this could exist (we could think of non-routine manual workers), but if we did allow for the possibility that low-skilled workers gain and high-skilled lose, this would require different examples and types of jobs (since it would not make sense in the scenario of a computer-based productivity improving innovation). That possibility would introduce more complexity and a much higher number of scenarios, which would raise power concerns. Finally, while we allow for price change to be 0 and for the change in the number of high-skilled workers to be 0, we do not allow the change in the number of low-skilled workers to be 0 since it would imply a decision with no trade-offs.

<sup>9</sup>For simplicity we keep the same prices across all four countries, while only changing the currencies. We do however adjust the wages based on the average wage for the occupation in question for each country according to the source talent.com. The wages shown in the table apply to the US and Canada. In the Australian case, the average wage for assembly workers is AUS\$50,000 (and the two scenarios after innovation are as a result, AUS\$30,000 or AUS\$40,000), while the average wage for data scientists is around AUS\$100,000; in the UK we adjust wages for both occupations. The average wage for data scientists is £60,000 (and the two scenarios after innovation are as a result, £75,000 or £90,000), while for assembly workers the average wage is £20,000 (and the two scenarios after innovation are as a result, £13,000 or £17,000).



Table 1: Table showing the effects of the introduction of a productivity-improving innovation. Square brackets contain the pre-specified set of possible values of attributes. Words in *italics* do not appear in the conjoint experiment, and are included here for clarity.

	Before Innovation	After Innovation
Firm	[Electronics; Aviation; Auto; Pharmaceutical]	[Electronics; Aviation; Auto; Pharmaceutical]
Price of [Smartphone; Plane; Car; Vaccine]:	<i>Smartphone</i> : \$600; <i>Plane</i> : \$100M; <i>Car</i> : \$25,000; <i>Vaccine</i> : \$25	<i>Smartphone</i> : [\$600; \$480; \$300]; <i>Plane</i> : [\$100M; \$80M; \$50M]; <i>Car</i> : [\$25,000; \$20,000; \$12,500]; <i>Vaccine</i> : [\$25; \$20; \$12.50]
Number of High Skilled Workers	200	[200, 250, 350]
Wage of High Skilled Workers	\$100,000	[\$125,000; \$150,000]
Number of Low Skilled Workers	200	[150, 50]
Wage of Low Skilled Workers	\$30,000	[\$20,000; \$25,000]

disagreed with the following statements, using a five-point scale from strongly agree to strongly disagree (with a sixth “don’t know” option):

- The company’s decision to introduce the new technology is fair.
- If you were the CEO of the company you would make the same decision to introduce the new technology.

After the last table in the conjoint, respondents were asked how much they agreed or disagreed with each of the policies listed above under the news article treatment section. This is meant to capture the aggregate effect of the high information treatment, since we do not ask the policy questions after each conjoint table.

## 4.2 Other measures

Aside from a battery of socio-demographic questions (gender, state of residence and postal code, education, income, marital status, employment status, citizenship and race, socioeconomic status, and political identification) the survey also includes a series of questions aimed to measure objective and subjective knowledge of automation. We are able to build an index of objective knowledge for individuals in the high information conjoint treatment group. As mentioned

earlier, after the first table, respondents in the conjoint were asked two questions to measure whether they understood the effects of innovation. The resulting index ranges from 0 to 2, to reflect the number of correct answers. Since these questions may be measuring attention rather than actual knowledge, we also validate our findings with an alternative variable that measures self-reported subjective knowledge of automation, which may make a respondent more likely to know that automation entails both costs and benefits and in turn may make them more sensitive to changes in these numbers. Individuals were asked:

- “How much would you say you understand automation” [I know nothing about it; I have heard the concept, but I don’t understand it well; I am familiar with the concept; I have a basic understanding; I have a good understanding; I am an expert]

The answers are recoded into three categories:

- “Low knowledge” [I know nothing about it; I have heard the concept, but I don’t understand it well];
- “Medium knowledge” [I am familiar with the concept; I have a basic understanding];
- “High knowledge” [I have a good understanding; I am an expert].

## 5 Results

### 5.1 Attitudes towards automation

The between-group analyses for the dependent variables *fairness* and *CEO* are shown in figures 1 and 2.<sup>10</sup> Figure 1 shows that people in both the low and high information treatment groups are more likely to think automation is fair than people in the news treatment group, who only see the costs of automation. On a scale from 1 (Strongly Disagree) to 5 (Strongly Agree) the predicted value

---

<sup>10</sup>See table A1 in the Appendix for tabular results.

of support for automation is 3.23 [95% CI: 3.17, 3.29] for people in the news information group, 3.68 [95% CI: 3.63, 3.73] for people in the low information group, and 3.37 [95% CI: 3.35, 3.40] for people in the high information conjoint group.

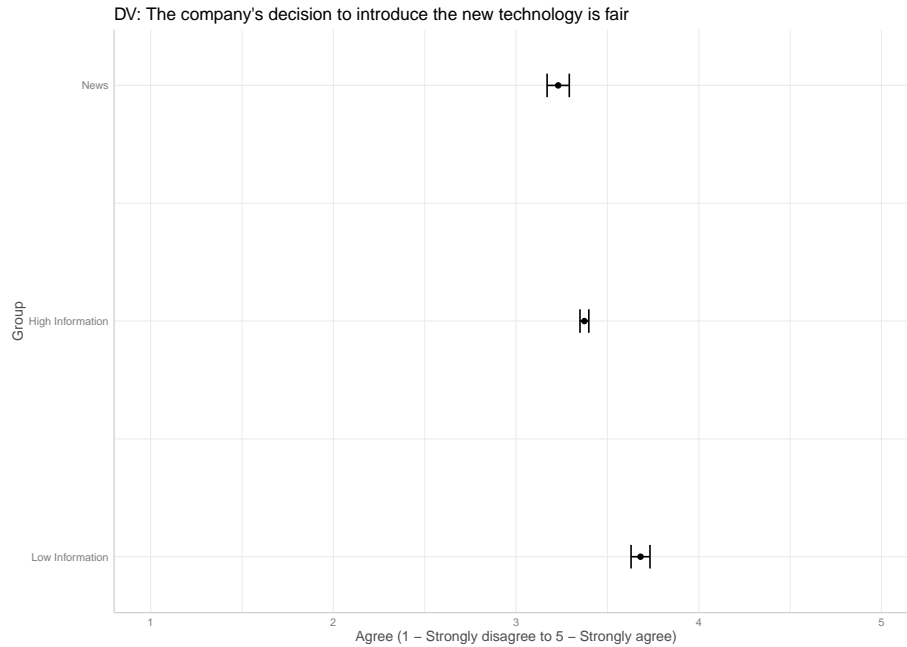


Figure 1: Coefficients from OLS regression on fairness with 95% confidence intervals.

Contrary to what we expected, people in the low information treatment group, who only see vague information about the trade-offs of automation, have more favorable views of automation than individuals in the high information group, who see precise numbers of winners and losers. This suggests that when we do not provide the exact numbers, respondents may imagine a more positive impact of automation than the average scenario in the high information conjoint treatment group. To contextualize this, on average across the four random vignettes in the conjoint treatment group, people saw: a fall in prices of about 21%, an increase of 67 high skilled workers (from 200), and a decrease of 100 low skilled workers (from 200), for a net loss of 33 workers. The average wage for high

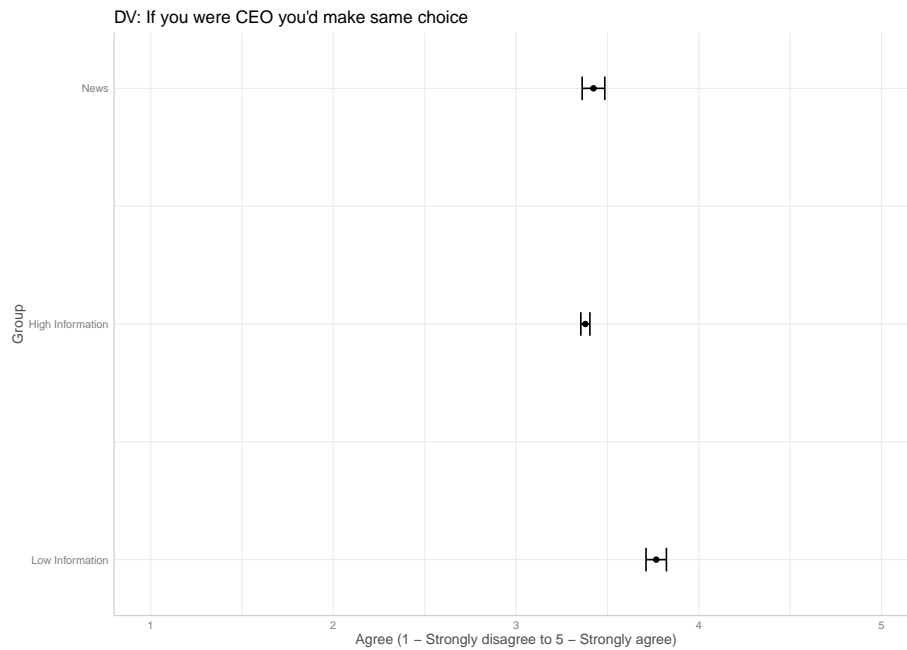


Figure 2: Coefficients from OLS regression on CEO with 95% confidence intervals.

skilled workers after the innovation was \$138,000 (from \$100,000), an increase of 38%, and for low skilled workers it was \$22,500 (from \$30,000), a decrease of 25%. The average scenario in the high information conjoint treatment may show greater costs or lower benefits than people would have expected in the vague trade-off treatment, thereby yielding lower average support for automation when compared to the latter condition.

Figure 2 tells a similar story for the dependent variable asking individuals whether they would have made a similar decision to automate if they were the CEO of the hypothetical company. However in this case the effects of the high information conjoint treatment and of the news treatment are not significantly different from each other.

## 5.2 Support for policies addressing job loss from automation

Next, we turn to the effects of the three treatments on support for different policy responses to automation. We hypothesized that in response to automation-driven job losses, people in the high information conjoint condition would be more likely to support redistributive policies than policies protecting jobs, compared to the other groups. We expected individuals in the costs-only news condition to be more supportive of policies protecting jobs than those involving labour market interventions, compared to individuals in the other groups. Figure 3 shows that this is not the case, as there appear to be no differences in support for different policies across the treatment groups. Overall, the favored policy response across all groups is retraining workers, which is also the most efficient of these policies, according to the literature.<sup>11</sup> Conversely, the policies with the lowest support across all three groups are those restricting skilled migration and imposing trade restrictions.

Consistently with these results, the between-group regression analyses of support for the eight individual policy responses show no effect of the high and low information treatments relative to the news article condition (see figures A1 to A8 and table A2 in the Appendix for tabular results). This suggests that regardless of the varying costs and benefits of automation, preferences for different policies are relatively sticky.

## 5.3 Conjoint analysis of attitudes towards automation

Next we turn to the conjoint analysis to investigate which attributes causally increase or decrease support for automation. Results for the dependent variable *fairness* are shown in the main text, while those for *CEO* are reported in the Appendix (see figures A9 to A12), since they are similar. Figure 4 shows

---

<sup>11</sup>Jaimovich et al. (2020) develop a macroeconomic model that examines the net effects of a menu of potential policy responses to the labour market effects of automation and find that the best policies that tend to help low-wage workers and promote economic growth are those that help displaced workers move to new jobs, such as retraining and unemployment insurance.

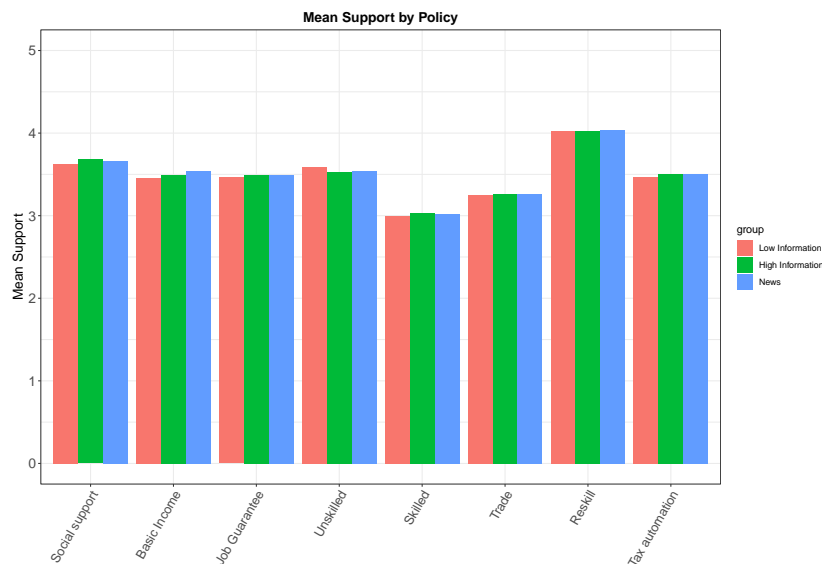


Figure 3: Mean support by policy by treatment group. The respondent indicates how much they agree that the government should implement each of the following policies, from 1 strongly disagree to 5 strongly agree.

the marginal means of the conjoint analysis, which represent the mean outcome across all appearances of a particular conjoint feature level, averaging across all other features. These suggest that people are most sensitive to price changes. Average marginal component effects from figure 5 confirm these results. Averaging across all other features, when prices decrease by 50% the mean support for automation is 3.52 [95% CI: 3.48, 3.55], while it is only 3.19 [95% CI: 3.15, 3.22] when there is no price change.

Similarly, when fewer (more) workers are lost (gained), respondents are more supportive of automation than when more (fewer) workers are lost (gained). They are also sensitive to varying declines in the wages of workers affected by automation, but not to wage increases, which suggests that gains and losses in wages are perceived differently. This is consistent with prospect theory, which predicts that people are more sensitive to economic losses than economic gains (Kahnemann & Tversky, 1979). Finally, respondents do not appear to be sensitive to the type of firm or product in question, whether it is an auto company




Figure 4: Marginal means for DV fairness

producing cars, an aviation firm producing planes, an electronics company making smartphones, or a pharmaceutical company producing vaccines.

Due to power issues we cannot estimate the marginal means for the best-case scenario compared to the worst-case scenario, in which we set each profile feature to its most optimistic or most pessimistic value, but we can do so for up to three features without sacrificing too much sample size. When we set the price decrease to 50%, the total number of high-skilled jobs post-innovation to 350 (the largest increase) and low-skilled jobs post-innovation to 150 (the lowest decrease), averaging across all other features, support for automation is 3.60 [95% CI: 3.54, 3.67]. Conversely, when we set the price decrease to 0%, the number of high-skilled jobs post-innovation to 200 (no change) and low-skilled jobs post-innovation to 50 (the largest decrease), averaging across all other features, support for automation is 3.15 [95% CI: 3.08, 3.23]. The point estimate of the best-case scenario is essentially equivalent to that of the




Figure 5: Average marginal component effects for DV fairness

low information group, suggesting again that the net benefits that respondents imagine in the low information treatment group are higher than those of the average conjoint scenario.

In figures 6 and 7 we test whether the observed effects differ by respondents' objective and subjective knowledge.<sup>12</sup> People who correctly estimated the costs and benefits of automation (who got two out of two correct answers), are more sensitive to price changes and to changes in the number of employed high-skilled workers than those who didn't compute those numbers correctly (0 out of 2 correct), as figure 6 illustrates.<sup>13</sup> Figure 7 shows that results are very similar for self-reported, as opposed to objective knowledge: People with higher self-reported knowledge of automation are more sensitive to price changes and to changes in the number of employed high-skilled workers than those with low

<sup>12</sup>We present results by high and low subjective and objective knowledge in the main text, since this is our comparison of interest. We present the full results, which include medium knowledge, in the Appendix (See figures A13 and A14).

<sup>13</sup>These models hold education constant since it may be a proxy for math skills.






Figure 6: Average marginal component effects for DV fairness by high and low knowledge

self-reported knowledge.

## 6 Conclusion

Technological change has arguably been the most important source of structural labour market transformations in advanced economies. However, its impacts on labour markets, and even more so on political and policy preferences, have only recently been gaining scholarly attention. Although projections of the future impact of technological change vary significantly, predicting modest to large effects, automation and AI will drastically transform the nature of work. The distributive consequences arising from the introduction of these new technologies have sparked academic debates on the political consequences of such transformations. Losers from technological change could possibly revolt and pose a threat to democratic stability if governments will fail to compensate them. But

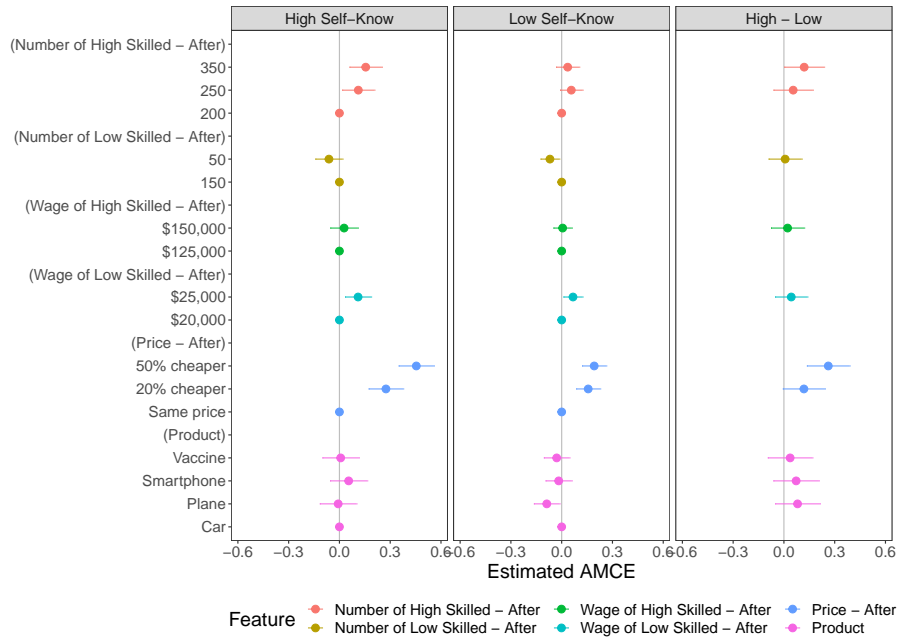


Figure 7: Average marginal component effects for DV fairness by high and low self-reported knowledge

the negative consequences of technological change are only one side of the story. These transformations will also create a larger group of technology beneficiaries. However, despite the dominant role of technological change in reshaping labour markets, most political parties have not claimed ownership of the issue or discussed specific policies in response to it; instead they have found it more politically advantageous to campaign against other issues, such as trade and immigration, which offer clearer culprits and solutions. Furthermore, much of the literature has mostly focused on the negative consequences of technological change (e.g., job loss), failing to discuss its potential for job gains and lower prices. For this reason, voters may fail to fully understand the consequences of technological change in the workplace and the implications for them and for others. We answer a series of questions in this paper. How do people perceive the trade-offs of automation? What policies in response to it do they favor and does support vary based on different net welfare effects of automation?

Overall, our findings suggest that relative to fear-mongering messages that only convey the costs of automation, support for automation increases significantly when generic, and to a lesser extent specific, information on its trade-offs is provided. People are particularly sensitive to price changes, and show the largest increases in support for automation when price declines are higher. Furthermore, the policy in response to automation with the highest support among respondents is retraining individuals negatively affected by technological change. Our results suggest that technological change may not necessarily result in the political backlash anticipated by many. At the same time, people are likely to be sensitive to information about the trade-offs of automation. According to our findings, they are less likely to approve of automation when it does not result in price decreases, and to a lesser extent, when it leads to lower job gains or larger job losses. As a result, how political parties choose to frame automation—for instance, whether they emphasize its price-reducing potential or its broader implications for labour markets—will affect voters’ perceptions of the issue and the degree to which they see it as a political problem or an economic opportunity.

## References

- Abbott, R., & Bogenschneider, B. (2018). Should robots pay taxes? tax policy in the age of automation. *Harvard Law and Policy Review*, *12*, 145–175.
- Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2020). What Are the Labor and Product Market Effects of Automation? New Evidence from France. *Working Paper*. [https://scholar.harvard.edu/files/aghion/files/what\\_are\\_the\\_labor\\_and\\_product\\_market\\_effects\\_of\\_automation\\_jan2020.pdf](https://scholar.harvard.edu/files/aghion/files/what_are_the_labor_and_product_market_effects_of_automation_jan2020.pdf)
- Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2021). The Direct and Indirect Effects of Automation on Employment: A Survey of the Recent Literature. *Working Paper*. [https://scholar.harvard.edu/files/aghion/files/direct\\_and\\_indirect\\_effects\\_of\\_automation.pdf](https://scholar.harvard.edu/files/aghion/files/direct_and_indirect_effects_of_automation.pdf)
- Anelli, M., Colantone, I., & Stanig, P. (2019). We Were the Robots: Automation and Voting Behavior in Western Europe. *IZA Discussion Paper Series. No. 12485*. <https://www.iza.org/publications/dp/12485/we-were-the-robots-automation-and-voting-behavior-in-western-europe>
- Autor, D. H., & Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, *103*(5), 1553–97.
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, *118*(4), 1279–1333.
- Autor, D. H., Mindell, D. A., & Reynolds, E. B. (2020). *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines*. The MIT Press.
- Bessen, J. (2019). Automation and jobs: when technology boosts employment. *Economic Policy*, *34*(100), 589–626.
- Borwein, S., Bonikowski, B., Loewen, P. J., Lee-Whiting, B., & Magistro, B. (2021). Automation, local and global economic dynamics, populism and nationalism. *Working Paper*.

- Borwein, S., Bonikowski, B., Loewen, P. J., Lee-Whiting, B., & Magistro, B. (2022). Perceived automation threat, populism, and vote choice: Evidence from 15 European democracies. *Working Paper*.
- Busemeyer, M. R., Gandenberger, M., Knotz, C., & Tober, T. (2022). Preferred policy responses to technological change: Survey evidence from OECD countries. *Socio-Economic Review*. <https://doi.org/10.1093/ser/mwac015>
- Busemeyer, M. R., & Sahm, A. H. (2021). Social investment, redistribution or basic income? exploring the association between automation risk and welfare state attitudes in europe. *Journal of Social Policy*, 1–20.
- Centre for the New Economy and Society. (2018). *The Future of Jobs Report 2018 Insight* (tech. rep.). [http://www3.weforum.org/docs/WEF\\_Future\\_of\\_Jobs\\_2018.pdf](http://www3.weforum.org/docs/WEF_Future_of_Jobs_2018.pdf)
- Colantone, I., & Stanig, P. (2018). The trade origins of economic nationalism: Import competition and voting behavior in Western Europe. *American Journal of Political Science*, 62(4), 936–953.
- Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6), 3104–3153.
- De La Rica, S., & Gortazar, L. (2016). Differences in Job De-Routinization in OECD Countries: Evidence from PIAAC. *IZA DP No. 9736*.
- Dermont, C., & Weisstanner, D. (2020). Automation and the future of the welfare state: Basic income as a response to technological change? *Political Research Exchange*, 2(1), 1–11.
- Di Tella, R., & Rodrik, D. (2020). Labour Market Shocks and the Demand for Trade Protection: Evidence from Online Surveys. *The Economic Journal*, 130(628), 1008–1030.
- Frey, C. B., Berger, T., & Chen, C. (2018). Political machinery: Did robots swing the 2016 US presidential election? *Oxford Review of Economic Policy*, 34(3), 418–442.

- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280.
- Gallego, A., Kuo, A., Manzano, D., & Fernández-Albertos, J. (2021). Technological Risk and Policy Preferences. *Comparative Political Studies*, *55*(1), 60–92.
- Gallego, A., & Kurer, T. (2022). Automation, Digitalization, and Artificial Intelligence in the Workplace: Implications for Political Behavior. *Annual Review of Political Science*, *25*(1).
- Gallego, A., Kurer, T., & Scholl, N. B. (2022). Neither Left-Behind nor Superstar: Ordinary Winners of Digitalization at the Ballot Box. *Journal of Politics*, *84*(1), 418–436.
- Goos, M., Manning, A., & Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, *104*(8), 2509–26.
- Guarascio, D., & Sacchi, S. (2021). Technology, risk and social policy. an empirical investigation. *GLO Discussion Paper*, (833).
- Humlum, A. (2019). Robot Adoption and Labor Market Dynamics. *Working paper*. <https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/5dcf78576d59eb44eac86f63/1573877848584/humlumJMP.pdf>
- Im, Z. J. (2021). Automation risk and support for welfare policies: How does the threat of unemployment affect demanding active labour market policy support? *Journal of International and Comparative Social Policy*, *37*(1), 76–91.
- Im, Z. J., Mayer, N., Palier, B., & Rovny, J. (2019). The “losers of automation”: A reservoir of votes for the radical right? *Research & Politics*, *6*(1).
- Iversen, T., & Soskice, D. (2001). An asset theory of social policy preferences. *American Political Science Review*, *95*(4), 875–893.
- Jaimovich, N., Saporta-Eksten, I., Siu, H., & Yedid-Levi, Y. (2020). The Macroeconomics of Automation: Data, Theory, and Policy Analysis. *Economic Studies at Brookings*.

- Jeffrey, K. (2021). Automation and the future of work: How rhetoric shapes the response in policy preferences. *Journal of Economic Behavior and Organization*, (192), 417–433.
- Jeffrey, K., & Matakos, K. (2021). Economic Anxiety, Anti-Elite Sentiment, and Demand for Redistribution. *QPE Working Paper*, (29). <http://dx.doi.org/10.2139/ssrn.3791279>
- Kahnemann, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292.
- Kurer, T. (2020). The declining middle: occupational change, social status, and the populist right. *Comparative Political Studies*, 53(10-11), 1798–835.
- Kurer, T., & Gallego, A. (2019). Distributional consequences of technological change: Worker-level evidence. *Research & Politics*, 6(1).
- Magistro, B. (2022). Party cues or policy information? The differential influence of financial and economic literacy on economic policy preferences. *Journal of Public Policy*, 1–24.
- Magistro, B., Lee-Whiting, B., Loewen, P., Borwein, S., & Bonikowski, B. (2022). Automation and demand for policies addressing job loss: the effects of information about trade-offs. *Pre-Analysis Plan at OSF*. <https://doi.org/10.17605/OSF.IO/A8X52>
- Menaldo, V. (Forthcoming). *Twilight of the Third Industrial Revolution and Dawn of the Fourth*. Manuscript in Progress.
- Mutz, D. (2021). (Mis)Attributing the Causes of American Job Loss: The Consequences of Getting it Wrong. *Public Opinion Quarterly*, 85(1), 101–122.
- Rodrik, D. (2018). Populism and the economics of globalization. *Journal of International Business Policy*, (1), 12–33.
- Sacchi, S., Guarascio, D., & Vannutelli, S. (2020). Risk of technological unemployment and support for redistributive policies. In R. Careja, P. Emmenegger, & N. Giger (Eds.), *The European Social Model under Pressure* (pp. 277–295). Springer VS, Wiesbaden.

- Thewissen, S., & Rueda, D. (2019). Automation and the welfare state: Technological change as a determinant of redistribution preferences. *Comparative Political Studies*, 52(2), 171–208.
- Werfel, S., Witko, C., & Heinrich, T. (2022). Public Support for Assistance for Workers Displaced by Technology. *Research & Politics*, 9(2), 1–7.
- Wu, N. (2021). Misattributed blame? Attitudes toward globalization in the age of automation. *Political Science Research and Methods*, 1–18.



## Appendix: For Online Publication

Table A1: Regression analyses on DVs fairness and CEO. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>	
	Fairness	Would do same if CEO
	(1)	(2)
High Information (ref. News)	0.143 (0.033)	-0.044 (0.034)
Low Information	0.451 (0.041)	0.343 (0.043)
Constant	3.231 (0.031)	3.424 (0.032)

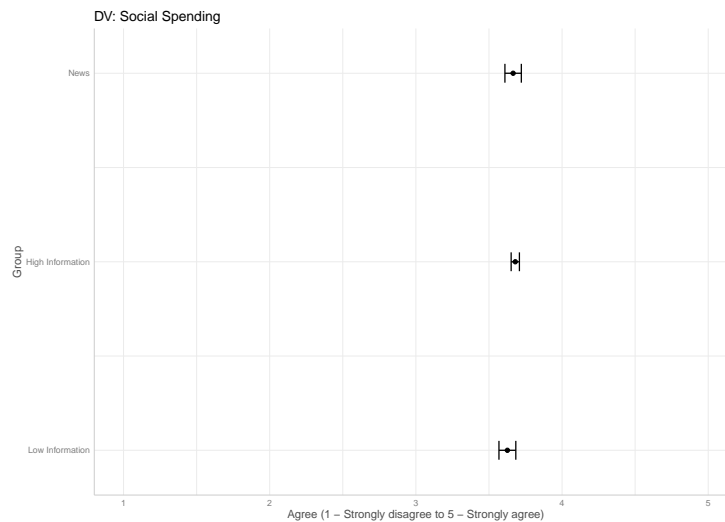


Figure A1: Coefficients from OLS regression on social spending with 95% confidence intervals.

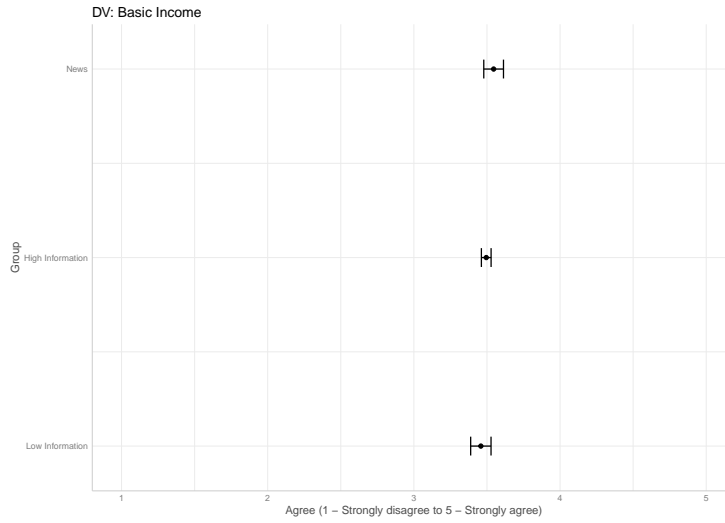


Figure A2: Coefficients from OLS regression on basic income with 95% confidence intervals.

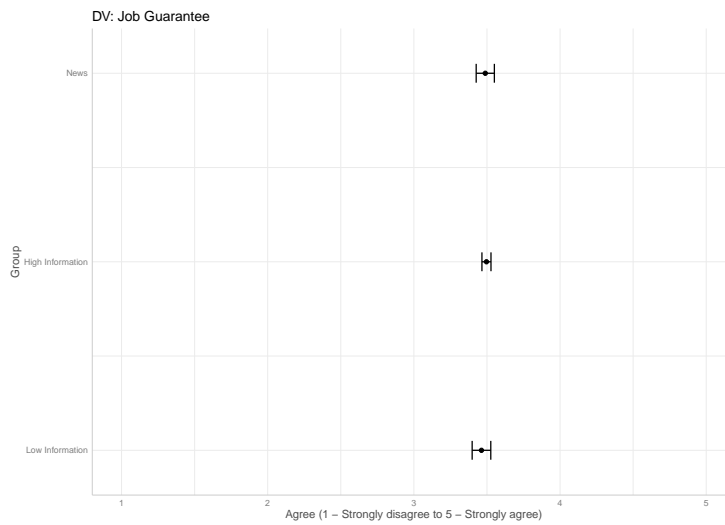


Figure A3: Coefficients from OLS regression on job guarantee with 95% confidence intervals.

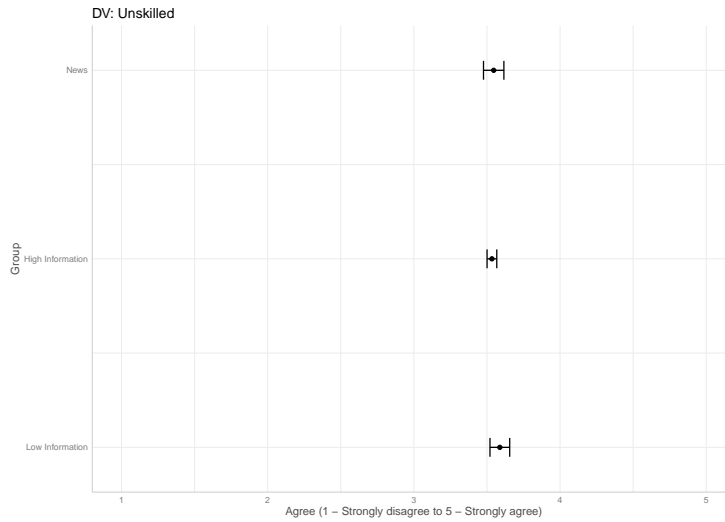


Figure A4: Coefficients from OLS regression on restricting unskilled migration with 95% confidence intervals.




Figure A5: Coefficients from OLS regression on restricting skilled migration with 95% confidence intervals.

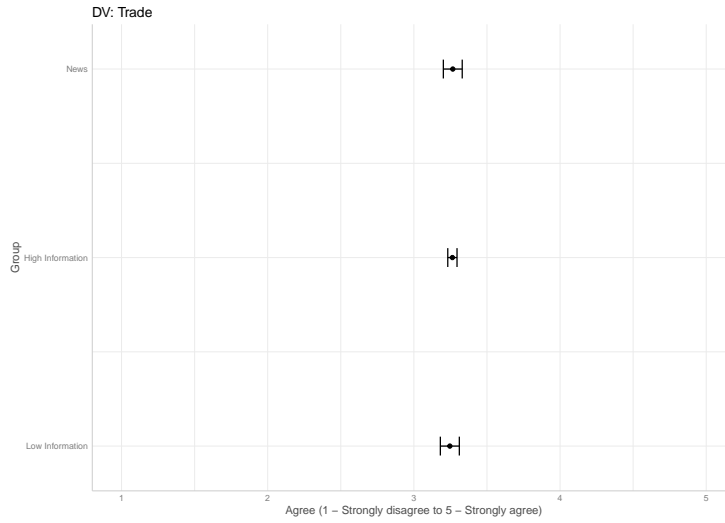


Figure A6: Coefficients from OLS regression on restricting trade with 95% confidence intervals.




Figure A7: Coefficients from OLS regression on retraining workers with 95% confidence intervals.




Figure A8: Coefficients from OLS regression on taxing automation with 95% confidence intervals.




Figure A9: Marginal means for DV CEO




Figure A10: Average marginal component effects for DV CEO




Figure A11: Average marginal component effects for DV CEO by knowledge




Figure A12: Average marginal component effects for DV CEO by self-reported knowledge






Figure A13: Average marginal component effects for DV fairness by knowledge




Figure A14: Average marginal component effects for DV fairness by self-reported knowledge

Table A2: Regression analyses of treatment group on policy support. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>							
	Social Spending	Basic Income	Job Guarantee	Unskilled	Skilled	Trade	Reskilling	Tax Automation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Information (ref. News)	0.015 (0.032)	-0.051 (0.038)	0.008 (0.035)	-0.012 (0.039)	0.014 (0.041)	-0.003 (0.037)	-0.006 (0.031)	-0.005 (0.039)
Low Information	-0.039 (0.041)	-0.087 (0.049)	-0.026 (0.045)	0.042 (0.050)	-0.017 (0.052)	-0.020 (0.047)	-0.005 (0.039)	-0.043 (0.049)
Constant	3.665 (0.029)	3.546 (0.034)	3.488 (0.032)	3.546 (0.036)	3.017 (0.037)	3.266 (0.033)	4.034 (0.028)	3.506 (0.035)

## Results by country

In this analysis we are dealing with relatively similar countries: all are liberal market economies, in which the same number of respondents received the same information treatments, and respondents were made representative in a similar way applying age, gender, and regional quotas. However, it is possible that a pooled analysis may hide some significant cross-country variation. Hence, we also run the analyses by country and report all the results below. Findings are not significantly or substantively different, with a few exceptions. Looking at the effects of the low information treatment on attitudes toward automation relative to the news article condition, we see that there are no significant differences across countries: people in the low information treatment group have more positive attitudes toward automation than those in the news article group. However, there is some variation when it comes to the high information treatment group. In Australia and the UK, people in the high information group are no more or less likely to think the company's decision to automate is fair than people in the news article treatment group. Conversely, this effect is positive and significant in Canada and the US: individuals in the high information group are more likely to think that the company's decision to automate is fair than those in the news article group. When it comes to what the respondent would do if they were the CEO of the company, there's slightly more variation. In Australia and the UK respondents in the high information group are less likely to think they would make the same decision to automate if they were CEO of the company compared to people in the news treatment group, while high information respondents in the US are more likely to do so, relative to the news condition. Finally, the effect of high information in Canada is not significantly different from that of the news article condition for the CEO dependent variable.

When it comes to policy preferences in response to job loss due to automation, respondents in the US overall display lower support for any policy response relative to the other three countries, except for restrictions on skilled migration and on trade, which they favor relatively more. Furthermore, while at the aggre-

gate level there are no significant treatment effects on policy preferences, there is some variation at the country-level. In the US, people in the low information treatment group are less supportive of social spending in response to job loss due to automation than respondents in the news article group. Similarly, in Canada, people in the low information treatment group are less supportive of trade restrictions than those in the news article group.

The conjoint analysis allows us to compare the effects of different attributes within the high information treatment condition. Overall, in Australia, Canada, and the UK respondents appear to be sensitive to prices the most, but also to changes in wages and in the number of jobs, consistently with the aggregate results. Respondents in the US are only sensitive to price changes, and not to any number of employees or wage changes. This may suggest that in the American context prices are more salient than other attributes when labour market changes are a result of automation, rather than other forces, such as offshoring. Finally, when looking at the conjoint analyses by knowledge and by country, while the effects are in the same direction as those at the aggregate level, whereby more knowledgeable people are more sensitive to price changes, in most cases the coefficients fail to reach statistical significance: this may be simply due to power issues at the individual country level. Overall, the country-by-country analysis shows that, with a few exceptions, respondents in the four different countries do not display different attitudes towards automation or demand different policies in response to job loss when presented with the same type of information.

Table A3: Regression analyses on DV fairness by country. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>			
	Fairness			
	Australia	Canada	UK	US
	(1)	(2)	(3)	(4)
High Information (ref. News)	-0.021 (0.063)	0.225 (0.067)	0.042 (0.065)	0.323 (0.069)
Low Information	0.331 (0.078)	0.595 (0.081)	0.349 (0.079)	0.522 (0.087)
Constant	3.365 (0.058)	3.168 (0.063)	3.255 (0.061)	3.142 (0.065)

Table A4: Regression analyses on DV CEO by country. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>			
	Would do the same if CEO			
	Australia	Canada	UK	US
	(1)	(2)	(3)	(4)
High Information (ref. News)	-0.155 (0.064)	-0.016 (0.068)	-0.153 (0.066)	0.146 (0.073)
Low Information	0.281 (0.081)	0.405 (0.083)	0.256 (0.083)	0.426 (0.092)
Constant	3.510 (0.059)	3.431 (0.063)	3.468 (0.061)	3.288 (0.068)

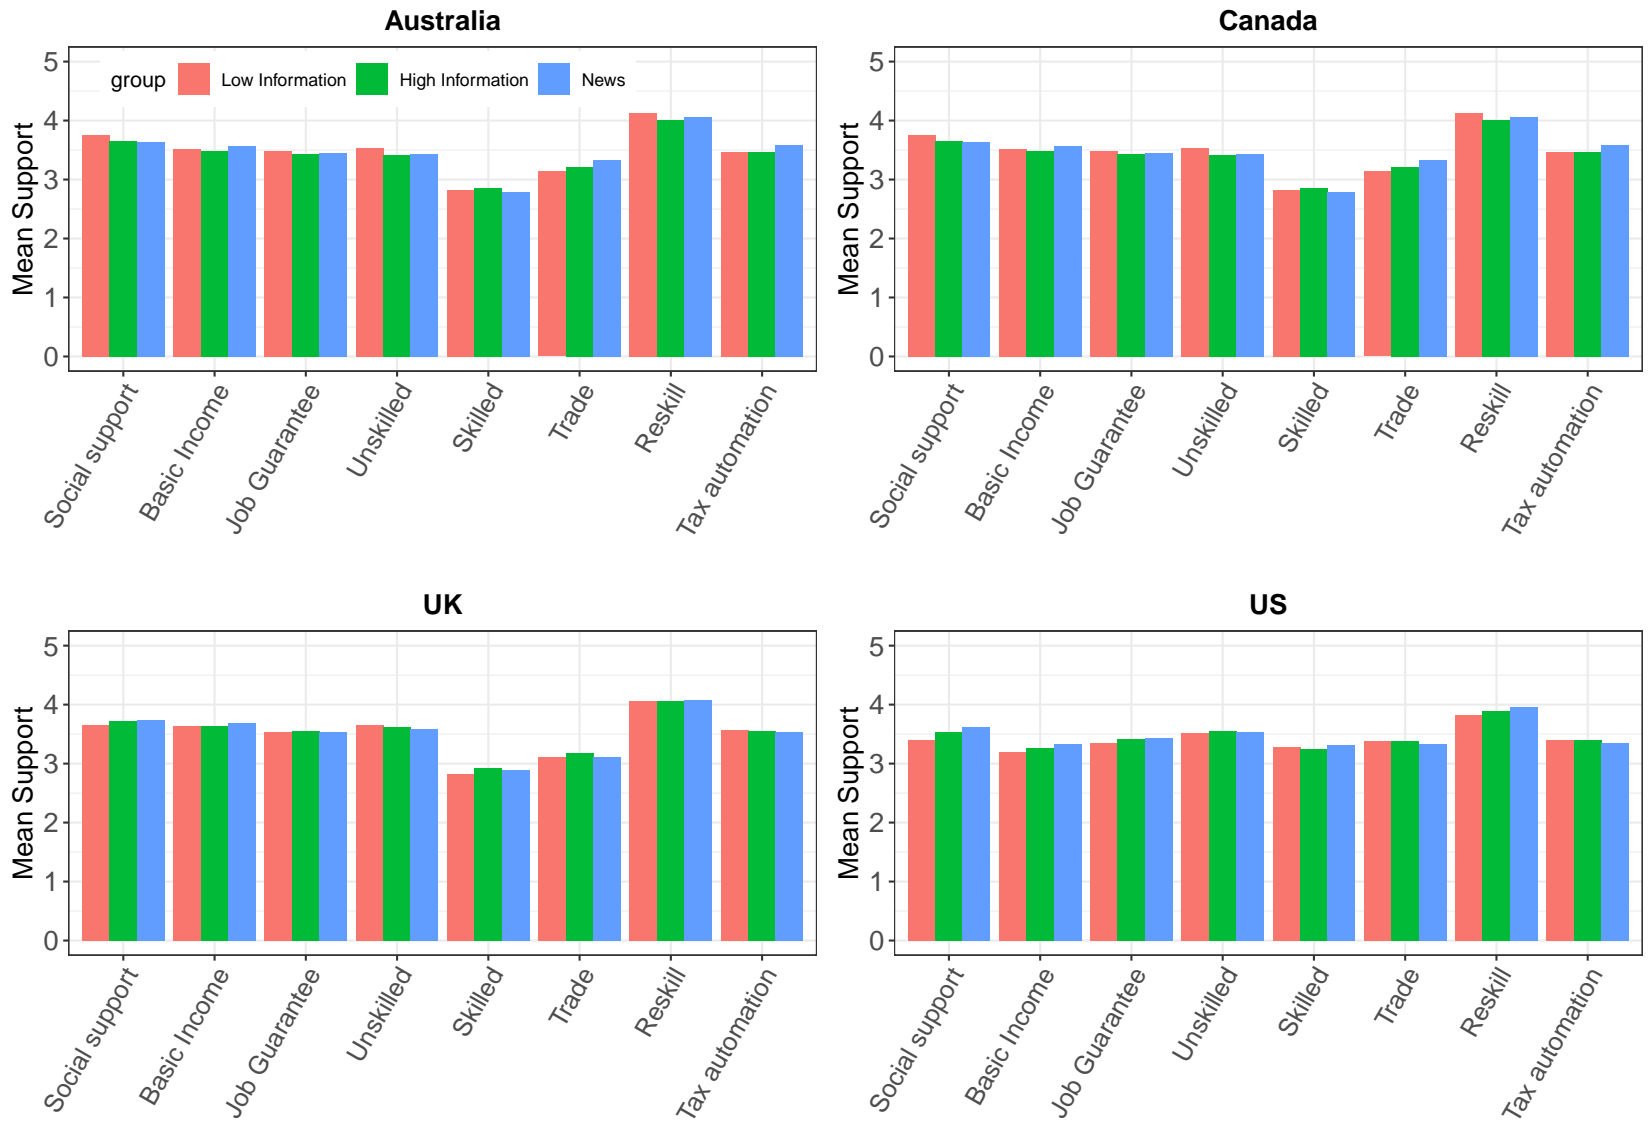


Figure A15: Mean Policy Support By Country and Treatment Group

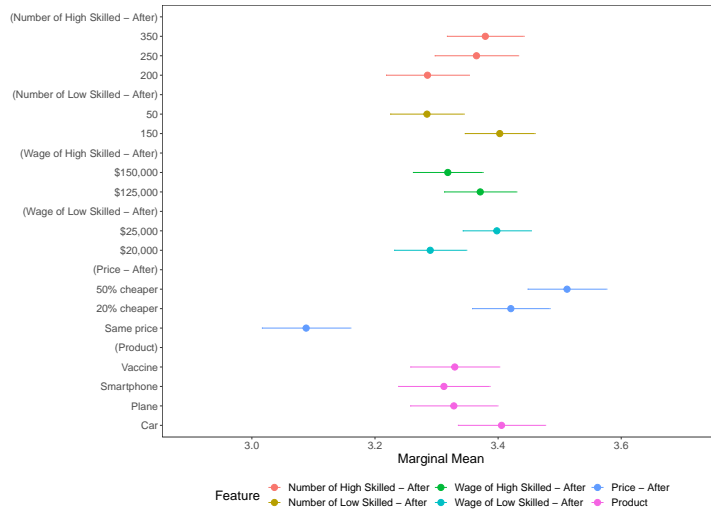


Figure A16: Marginal means for DV Fairness: Australia




Figure A17: Average marginal component effects for DV Fairness: Australia






Figure A18: Marginal means for DV CEO: Australia




Figure A19: Average marginal component effects for DV CEO: Australia




Figure A20: Marginal means for DV Fairness: Canada




Figure A21: Average marginal component effects for DV Fairness: Canada




Figure A22: Marginal means for DV CEO: Canada




Figure A23: Average marginal component effects for DV CEO: Canada




Figure A24: Marginal means for DV Fairness: UK




Figure A25: Average marginal component effects for DV Fairness: UK




Figure A26: Marginal means for DV CEO: UK




Figure A27: Average marginal component effects for DV CEO: UK




Figure A28: Marginal means for DV Fairness: US




Figure A29: Average marginal component effects for DV Fairness: US




Figure A30: Marginal means for DV CEO: US




Figure A31: Average marginal component effects for DV CEO: US

Table A5: Regression analyses of treatment group on policy support for Australia. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>							
	Social Spending	Basic Income	Job Guarantee	Unskilled	Skilled	Trade	Reskilling	Tax Automation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Information (ref. News)	0.121 (0.062)	-0.016 (0.073)	0.037 (0.070)	-0.097 (0.077)	-0.014 (0.080)	-0.022 (0.074)	0.076 (0.060)	0.025 (0.077)
Low Information	-0.004 (0.080)	-0.148 (0.094)	-0.059 (0.088)	0.005 (0.095)	0.016 (0.102)	0.073 (0.092)	0.032 (0.076)	-0.149 (0.098)
Constant	3.681 (0.056)	3.604 (0.065)	3.536 (0.063)	3.644 (0.069)	3.086 (0.072)	3.297 (0.066)	4.051 (0.054)	3.556 (0.070)



Table A6: Regression analyses of treatment group on policy support for Canada. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>							
	Social Spending	Basic Income	Job Guarantee	Unskilled	Skilled	Trade	Reskilling	Tax Automation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Information (ref. News)	0.028 (0.064)	-0.081 (0.078)	-0.011 (0.071)	-0.020 (0.079)	0.072 (0.079)	-0.113 (0.072)	-0.045 (0.060)	-0.115 (0.078)
Low Information	0.121 (0.081)	-0.048 (0.100)	0.039 (0.090)	0.100 (0.098)	0.029 (0.101)	-0.188 (0.090)	0.073 (0.073)	-0.113 (0.099)
Constant	3.630 (0.057)	3.561 (0.070)	3.443 (0.063)	3.429 (0.070)	2.791 (0.070)	3.325 (0.065)	4.056 (0.053)	3.581 (0.070)

Table A7: Regression analyses of treatment group on policy support for the UK. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>							
	Social Spending	Basic Income	Job Guarantee	Unskilled	Skilled	Trade	Reskilling	Tax Automation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Information (ref. News)	-0.010 (0.058)	-0.045 (0.068)	0.016 (0.065)	0.030 (0.077)	0.034 (0.078)	0.081 (0.071)	-0.002 (0.057)	0.030 (0.075)
Low Information	-0.093 (0.073)	-0.051 (0.089)	-0.008 (0.085)	0.062 (0.097)	-0.069 (0.097)	0.009 (0.091)	-0.010 (0.071)	0.036 (0.094)
Constant	3.740 (0.052)	3.686 (0.060)	3.540 (0.059)	3.585 (0.070)	2.897 (0.071)	3.103 (0.064)	4.070 (0.052)	3.530 (0.068)

Table A8: Regression analyses of treatment group on policy support for the US. Robust standard errors are given in parentheses.

	<i>Dependent variable:</i>							
	Social Spending	Basic Income	Job Guarantee	Unskilled	Skilled	Trade	Reskilling	Tax Automation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Information (ref. News)	-0.086 (0.072)	-0.067 (0.086)	-0.018 (0.077)	0.019 (0.082)	-0.057 (0.084)	0.045 (0.076)	-0.057 (0.068)	0.041 (0.080)
Low Information	-0.207 (0.094)	-0.135 (0.110)	-0.094 (0.100)	-0.019 (0.106)	-0.021 (0.109)	0.038 (0.098)	-0.135 (0.088)	0.050 (0.103)
Constant	3.613 (0.064)	3.331 (0.077)	3.439 (0.069)	3.537 (0.075)	3.308 (0.076)	3.337 (0.068)	3.957 (0.061)	3.356 (0.072)




Figure A32: Average marginal component effects for DV Fairness by Objective Knowledge: Australia




Figure A33: Average marginal component effects for DV CEO by Objective Knowledge: Australia




Figure A34: Average marginal component effects for DV Fairness by Self-Reported Subjective Knowledge: Australia




Figure A35: Average marginal component effects for DV CEO by Self-Reported Subjective: Australia




Figure A36: Average marginal component effects for DV Fairness by Objective Knowledge: Canada






Figure A37: Average marginal component effects for DV CEO by Objective Knowledge: Canada




Figure A38: Average marginal component effects for DV Fairness by Self-Reported Subjective Knowledge: Canada




Figure A39: Average marginal component effects for DV CEO by Self-Reported Subjective: Canada




Figure A40: Average marginal component effects for DV Fairness by Objective Knowledge: UK




Figure A41: Average marginal component effects for DV CEO by Objective Knowledge: UK




Figure A42: Average marginal component effects for DV Fairness by Self-Reported Subjective Knowledge: UK




Figure A43: Average marginal component effects for DV CEO by Self-Reported Subjective: UK




Figure A44: Average marginal component effects for DV Fairness by Objective Knowledge: US






Figure A45: Average marginal component effects for DV CEO by Objective Knowledge: US




Figure A46: Average marginal component effects for DV Fairness by Self-Reported Subjective Knowledge: US




Figure A47: Average marginal component effects for DV CEO by Self-Reported Subjective Knowledge: US