Automation and AI Research Project

Beatrice Magistro, Peter Loewen, Bart Bonikowski, Sophie Borwein, Blake Lee-Whiting

1 Research Questions

Modern economies are experiencing major changes, including labor dislocation, as a result of a new wave of automation and artificial intelligence (AI). This project attempts to answer multiple questions. The key question of the first part of the project is the following:

- Does providing information about the trade-offs of automation affect support for automation and, ultimately, support for different policies in response to it? Is there a gender gap with respect to such policy preferences?

The second part of the project investigates the following questions:

- According to citizens, when should governments use AI?
- What do citizens know about AI and algorithms?
- What are people's biggest concerns regarding AI?

2 Background

Recently, labor markets have been profoundly reshaped by technological change, in particular by automation and artificial intelligence (AI). While many argue that these developments will produce mass unemployment as new technologies replace labor, 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 one, enhancing productivity (Aghion et al., 2021; Aghion et al., 2020; Bessen, 2019). This latter approach emphasizes the direct productivity effect of automation: firms that automate become 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, Levy, and Murnane, 2003; Autor and Dorn, 2013). Some studies find evidence of a reallocation of workers between occupations (Humlum, 2019), with labor demand shifting from low-skilled to high-skilled workers. Kurer and Gallego (2019) zoom in 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. 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 labor 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 one looks at citizens' preferred policies to address technological change. Recent findings suggest that workers more at risk from job loss due to automation are more likely to support radical right parties in Europe (Anelli, Colantone, and Stanig, 2019; Im et al., 2019) and Donald Trump in the US (Frey, Berger, and Chen, 2018). This literature mostly focuses on left-behind voters, while neglecting the large majority of workers that benefits from innovation. One exception is Gallego, Kurer, and Scholl (2021), which studies how digitalization affects political preferences among the entire active labor force in the UK between 1997 and 2017 and provides evidence that digitalization was economically beneficial for workers with middle and high levels of education, while it produced

small negative effects for low education workers. Furthermore, growth in digitalization increased support for the incumbent party and voter turnout among the winners from digitalization.

However, automation is generally discussed in a negative light by the media as well as by the scholarly literature, as it is more often associated with job loss rather than job creation (Anelli, Colantone, and Stanig, 2019; Frey and Osborne, 2017; Im et al., 2019). This raises the question of whether most people are aware of the aggregate as well as the heterogeneous effects of automation. The majority of experimental scholarly works that look at automation (and similarly, trade) preferences uses a static framework, where information treatments only communicate the costs of these phenomena, stressing only the number of jobs that may be lost to automation. Building in particular on the work of Di Tella and Rodrik (2020) and Borwein et al. (2021), in the first part of our project we plan to expand the work on preferences for automation and conduct survey experiments that manipulate information to then infer respondents' support for different policy responses. Rather than just emphasizing costs, we will provide a broader picture of the potential effects of automation, designing different scenarios where we vary the numbers of winners and losers (in terms of final products' price changes and changes in the number of employed workers) to then infer people's preferences over policy responses to automation. We want to see how people weigh different trade-offs when evaluating automation and policy responses to it. Under what conditions are they more or less likely to favor automation and to support more efficient solutions, such as retraining workers or providing unemployment insurance for displaced workers, as opposed to protectionism?[1] (Jaimovich et al., 2020).

In answering this question, we anticipate that perceptions of automation, as well as preferred policy responses, will vary across respondents. Individuals who are at higher risk of automation (either objectively or subjectively so) may differ in their response to the treatments, since the weights they will place on the costs of automation will likely be higher. Furthermore, we know from the free trade literature that women tend to be more protectionist than men. We are very interested to investigate whether the same gender gap exists with respect to automation or not and

[1]Jaimovich et al. (2020) develop a macroeconomic model that serves as a theoretical laboratory to examine the net effects of a menu of potential policy responses to the labor 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.

whether women tend to support the same policies in response to automation as men once they see the information treatment. Finally, this effect is also likely to vary by knowledge. The effect size will be larger for people who correctly computed the costs and benefits of the new innovation.

The second part of the project will concern what role people perceive the government should have in AI. Furthermore, we will investigate how much citizens are familiar with AI and algorithms and what their fears around these are. In particular, today, many decisions that are made by human beings can actually be made by computer algorithms. While algorithmic decision-making brings several benefits, which potentially include an increase in accuracy and a reduction in human bias, others worry that these systems may just reinforce existing biases and disparities. When investigating the public's views and concerns on AI and algorithms it is also important to understand citizens' knowledge of these phenomena, as they will increasingly be asked to make decisions on these issues.

While these algorithms are already all around us, from book and movie recommendations, to news stories they think we may find relevant, they are increasingly also being used in many areas of government, including traffic management and customer service centers. These consist of computers analyzing large amounts of data, detecting statistical patterns, and developing models that can be used to make accurate predictions. Although these tools are likely to lead to an increase in efficiency and effectiveness in government policy, their use remains contentious. In particular, key concerns revolve around potentially perpetuating bias, creating negative effects on employment, and privacy, ethical, and transparency issues. We hence aim to investigate which government decisions people would feel comfortable to be made by computers rather than human beings, what their fears around the use of AI by governments are, and what they currently know about AI and algorithms. Previous findings (Carrasco et al., 2019) suggest that citizens are more supportive of government use of AI for tasks such as traffic optimization and transportation, or customer service activities than for more sensitive decisions, such as those involving the justice system or medical diagnoses. Furthermore, citizens were most concerned about ethical issues and the potentially negative effects of automation on employment. Trust in government, perceived corruption, and knowledge about AI and algorithms, are all likely to affect support for government for government use of AI, and we should hence expect heterogeneous effects based on these.

3 Research Design

3.1 Part 1

What happens after a firm introduces new computer-based productivity improving technology? Theoretically, we can think of a few different simple scenarios.

Let's assume we initially have a perfectly competitive market for the main good being produced by the firm in question. The new computer-based technology will make the firm more productive, shifting the supply curve out, this in turn will drive quality-adjusted prices down for consumers, and - provided that demand is elastic enough to prices - product demand will increase, which will result in net job growth. However, the expectation is that low-skilled and high-skilled workers in the firm will be affected heterogeneously by this 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 greatly. 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, as automation will reduce the demand for other 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 large enough that it would be more efficient to compensate the losers (retraining, unemployment insurance, or other options) than to stop innovation altogether.

In another scenario, if we assume now that demand is more inelastic, as prices go down as a result of innovation, demand will not increase enough to result in net job growth. Again, we could consider heterogeneous effects, where this time, only a minority of high-skilled workers benefits, while a majority of low-skilled ones loses.

Furthermore, if we relaxed 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.

Within all of these different scenarios we could hence vary the number of winners and losers from automation and see how people react to them.

Each person would be randomized into either a control scenario, a news article treatment, or a conjoint treatment. The conjoint will tell us what attributes causally increase or decrease support for automation on average when varied independently of the other attributes included in the design, and will allow us to compare the marginal effect of an information treatment on automation preferences to a newspaper article emphasizing the costs of automation only and to a control group that doesn't see numbers on who gains or loses.

3.1.1 Control group

The control vignette could be something along these lines, where individuals would hear about the new innovation and that it would involve trade-offs, but are not given any exact numbers:

> "A manufacturing firm in country X[vary country] decides to introduce a new computer-based productivity improving technology. As a result of this innovation, the cost to produce their main product will decrease, and its final price could also decrease. Furthermore, while some jobs will be gained, others will be lost."

After viewing the vignette, respondents are asked how much they agree or disagree with the following statements (from Borwein et al. 2021):

- If you were the CEO of the company you would make the same decision to introduce the new technology [strongly agree, agree, disagree, strongly disagree, don't know]
- The company's decision to introduce the new technology is fair [strongly agree, agree, disagree, strongly disagree, don't know]

Subsequently, they would be asked how much they agree or disagree with each of the following policies (from Borwein et al. 2021):

- Do you agree that the government should implement the following policy:
  - Expand social spending to support laid-off workers, and workers in similar positions. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Implement a basic income that gives every adult a set amount of money from government on a regular basis. [strongly agree, agree, disagree, strongly disagree, don't know]

- Pay to retain displaced workers and guarantee them jobs. [strongly agree, agree, disagree, strongly disagree, don't know]
- Reduce the number of unskilled immigrants entering the country for work. [strongly agree, agree, disagree, strongly disagree, don't know]
- Reduce the number of skilled immigrants entering the country for work. [strongly agree, agree, disagree, strongly disagree, don't know]
- Restrict international competition by increasing trade barriers on goods and services to the United States. [strongly agree, agree, disagree, strongly disagree, don't know]
- Fund programs to re-skill workers for new jobs. [strongly agree, agree, disagree, strongly disagree, don't know]
- Directly tax companies that replace workers with machines and robots. [strongly agree, agree, disagree, strongly disagree, don't know]
- Do nothing. [strongly agree, agree, disagree, strongly disagree, don't know]

3.1.2 News article treatment group

The news article treatment group vignette could be something along these lines, where individuals read, in sensational terms, about the costs of automation only, i.e., the job losses, but are not given information on the benefits:

> "Assembly and factory jobs are at risk at a manufacturing firm in country X[vary country], 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: "Many will become unemployed and the rest would have to accept lower wages," he added.

After viewing the vignette, respondents are asked how much they agree or disagree with the following statements (from Borwein et al. 2021):

- If you were the CEO of the company you would make the same decision to introduce the new technology [strongly agree, agree, disagree, strongly disagree, don't know]

- The company's decision to introduce the new technology is fair [strongly agree, agree, disagree, strongly disagree, don't know]

Subsequently, they would be asked how much they agree or disagree with each of the following policies (from Borwein et al. 2021):

- Do you agree that the government should implement the following policy:
  - Expand social spending to support laid-off workers, and workers in similar positions. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Implement a basic income that gives every adult a set amount of money from government on a regular basis. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Pay to retain displaced workers and guarantee them jobs. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Reduce the number of unskilled immigrants entering the country for work. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Reduce the number of skilled immigrants entering the country for work. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Restrict international competition by increasing trade barriers on goods and services to the United States. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Fund programs to re-skill workers for new jobs. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Directly tax companies that replace workers with machines and robots. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Do nothing. [strongly agree, agree, disagree, strongly disagree, don't know]

3.1.3 Information conjoint treatment

Individuals in the conjoint group will first see a pre-treatment vignette:

> "A manufacturing firm in country X[vary country] decides to introduce a new computer-based productivity improving technology. As a result of this innovation, the cost to

Table 1: Table showing the effects of the introduction of a productivity-improving innovation. In square brackets are the pre-specified set of possible values of attributes.

	Before Innovation	After Innovation
Firm	[Electronics, Aviation, Auto, Pharmaceutical]	[Electronics, Aviation, Auto, Pharmaceutical]
Price of Main Final Product [Cellphone, Plane, Car, Vaccine]	600	[600, 480, 300]
Number of High Skilled Workers	200,000	[200,000; 250,000; 350,000]
Wage of High Skilled Workers	$100,000	[$125,000; $150,000]
Number of Low Skilled Workers	200,000	[50,000; 150,000]
Wage of Low Skilled Workers	$50,000	[$30,000; $40,000]

produce their main product will decrease, and its final price could also decrease. This innovation could create new highly-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 be substituted by machines. Furthermore, the remaining low-skilled workers will also see a cut in their yearly pay. The table below shows the numbers in question."

Then, each individual would see a combination of charts, each with varying numbers for winners and losers[2].

[2]For simplicity, for now we keep one job per type (one high-skilled – AI, big data, and machine learning specialists – and one low-skilled – factor and assembly workers). We do not allow for a scenario where low-skilled workers actually win. Theoretically, this could exist (we could think of non-routine manual workers), but if we did allow for that possibility (that low-skilled workers gain and high-skilled lose) we would have to change the examples completely and the 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 don't allow the change in the number of low-skilled workers to be 0 since it would imply a decision with no trade-offs.

After the first table respondents are asked these questions to see whether they understood effects of innovation (multiple choice):

- 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 are asked how much they agree or disagree with the following statements (same as above for control group):

- If you were the CEO of the company you would make the same decision to introduce the new technology [strongly agree, agree, disagree, strongly disagree, don't know]
- The company's decision to introduce the new technology is fair [strongly agree, agree, disagree, strongly disagree, don't know]

After the last table in the conjoint, they would asked how much they agree or disagree with each of the following policies (so this would capture a compound effect of the information treatment, since we do not ask this one after each table):

- Do you agree that the government should implement the following policy:
  - Expand social spending to support laid-off workers, and workers in similar positions. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Implement a basic income that gives every adult a set amount of money from government on a regular basis. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Pay to retain displaced workers and guarantee them jobs. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Reduce the number of unskilled immigrants entering the country for work. [strongly agree, agree, disagree, strongly disagree, don't know]
  - Reduce the number of skilled immigrants entering the country for work. [strongly agree, agree, disagree, strongly disagree, don't know]

- Restrict international competition by increasing trade barriers on goods and services to the United States. [strongly agree, agree, disagree, strongly disagree, don't know]
- Fund programs to re-skill workers for new jobs. [strongly agree, agree, disagree, strongly disagree, don't know]
- Directly tax companies that replace workers with machines and robots. [strongly agree, agree, disagree, strongly disagree, don't know]
- Do nothing. [strongly agree, agree, disagree, strongly disagree, don't know]

3.2 Part 2

In the second part of the project we investigate what role people perceive the government should have in AI, how much citizens are familiar with AI and algorithms, and what their fears around these are.

In answering the question about how comfortable people are with decisions being made by a computer rather than a human being in government policy, we could use the questionnaire used by Carrasco et al. (2019), where individuals express their support for deployment of AI in government on a scale from 1 (totally disagree) to 7 (totally agree) on the following issues:

- Using real-time information to predict issues, optimize, and reroute traffic
- Predicting breakdowns and maintenance requirements for machinery and equipment
- Matching job seekers with available jobs
- Determining tax assessments and tax payable
- Providing virtual assistants to offer customer service and answer general inquiries
- Identifying potential fraud or noncompliance in administration of government services
- Assessing travelers for additional security screening
- Determining eligibility and amounts for welfare and social security entitlements
- Assessing medical images and making diagnoses (e.g. identifying cancer from scans)

- Determining eligibility for visas and immigration
- Making recommendations for medical treatment
- Making parole board decisions
- Determining innocence or guilt in a criminal trial.

In order to understand people's knowledge of AI and algorithms we could use both subjective and objective measures. As far as subjective measures are concerned, we could ask something along these lines:

- I understand what the term "artificial intelligence" means. [Strongly disagree, Disagree, Neither agree nor disagree, Agree, Strongly Agree]
- I know what an algorithm is in the context of computer science. [Strongly disagree, Disagree, Neither agree nor disagree, Agree, Strongly Agree]

To have an objective measure, one option is to use some questions from the Innovation, Science and Economic Development Canada (2021). The first one is:

Check all that apply from the following list of 11 proposed 'capabilities' that AI is able to perform at this time:

- Learn from data to increase understanding
- Perform video surveillance
- Interpret speech
- Play games
- Interpret images
- Replace humans doing dangerous tasks
- Help solve business problems
- Think logically

- Compose music
- Behave as humans do in social settings
- Feel emotion

Currently, AI is capable of performing all of the tasks except for the last two to some degree.
The second one is: Which of the following technologies use AI? (Check all that apply).

- Email spam filters
- Predictive search terms (i.e. predictions of what you are looking for based on popular search terms, etc.)
- Virtual assistant (i.e. Siri, Alexa, etc.)
- Online virtual assistant (i.e. Chat Bot, etc.)
- Recommender systems (i.e. online shopping, Netflix, etc.)
- Image search/recognition
- None use AI
- Unsure

All of these technologies are AI-enabled.
Finally, another option is to use the questions we had in Borwein et al. (2021):

- To what extent do you think technology drives progress? [Not at all, Somewhat, A lot]

We are now going to ask you some questions to help us learn about what types of information about AI and Automation are known to the public. Please answer these questions on your own without asking anyone or looking up the answers. Many people don't know the answers to these questions, so if you do not know the answer, please select "don't know".
Which of the following is an example of Artificial Intelligence?

- An electric toothbrush

- A car with an automatic transmission
- Siri, Alexa, and other voice-chat assistants
- A mechanical pencil
- None of the above
- I don't know

Which of the following is an example of Automation?

- Winemaker sampling wine from a wine barrel
- Cleaner mopping the floors of a high-school
- A person taking out cash from an ATM
- A chef preparing a cake using a spatula
- None of the above
- I don't know

What is a Turing Test? Please select one option.

- A complex set of questions intended to check the speed of an automotive process
- A test of a machine's ability to exhibit intelligent behavior indistinguishable from that of a human
- A test of a robot's ability to carefully lift complex objects within a controlled environment
- A type of test given to students in which a computer program automatically marks the student's answers
- None of the above
- I don't know

We could also ask two open-ended questions:

- When you think of automation, what comes to mind? [Open question]
- When you think of artificial intelligence (AI), what comes to mind? [Open question]

To analyze people's fears around government's deployment of AI we could also borrow in part from Carrasco et al. (2019). Specifically: "What concerns you the most about AI in government? (Check all that apply)"

- Violation of citizens' privacy
- Ethical implications (e.g. lack of fairness)
- Potential for bias and discrimination
- Lack of transparency in decision-making (e.g. lack of responsibility and accountability)
- Accuracy of results and analysis
- Capability of the public sector to use AI
- Impact of AI on jobs
- No concerns
- Other (Please write down)

References

Aghion, Philippe et al. (2020). "What Are the Labor and Product Market Effects of Automation? New Evidence from France". In: *Working Paper*. URL: https://scholar.harvard.edu/files/aghion/files/what_are_the_labor_and_product_market_effects_of_automation_jan2020.pdf.

— (2021). "The Direct and Indirect Effects of Automation on Employment: A Survey of the Recent Literature". In: *Working Paper*. URL: https://scholar.harvard.edu/files/aghion/files/direct_and_indirect_effects_of_automation.pdf.

Anelli, Massimo, Italo Colantone, and Piero Stanig (2019). "We Were the Robots: Automation and Voting Behavior in Western Europe". In: *IZA Discussion Paper Series. No. 12485*. URL: https://www.iza.org/publications/dp/12485/we-were-the-robots-automation-and-voting-behavior-in-western-europe.

Autor, David H. and David Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market". In: *American Economic Review* 103.5, pp. 1553–97. DOI: 10.1257/AER.103.5.1553.

Autor, David H., Frank Levy, and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration". In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333. DOI: 10.1162/003355303322552801.

Bessen, James (2019). "Automation and jobs: when technology boosts employment". In: *Economic Policy* 34.100, pp. 589–626. DOI: 10.1093/EPOLIC/EIAA001.

Carrasco, Miguel et al. (2019). "The Citizen's perspective on the use AI in government". In: *BCG Digital Government Benchmarking*. URL: https://www.bcg.com/publications/2019/citizen-perspective-use-artificial-intelligence-government-digital-benchmarking.

Centre for the New Economy and Society (2018). *The Future of Jobs Report 2018 Insight*. Tech. rep. URL: http://www3.weforum.org/docs/WEF_Future_of_Jobs_2018.pdf.

Di Tella, Rafael and Dani Rodrik (2020). "Labour Market Shocks and the Demand for Trade Protection: Evidence from Online Surveys". In: *The Economic Journal* 130.628, pp. 1008–1030. DOI: 10.1093/EJ/UEAA006.

Frey, Carl Benedikt, Thor Berger, and Chinchih Chen (2018). "Political machinery: Did robots swing the 2016 US presidential election?" In: *Oxford Review of Economic Policy* 34.3, pp. 418–442. DOI: 10.1093/OXREP/GRY007.

Frey, Carl Benedikt and Michael A. Osborne (2017). "The future of employment: How susceptible are jobs to computerisation?" In: *Technological Forecasting and Social Change* 114, pp. 254–280. DOI: 10.1016/J.TECHFORE.2016.08.019.

Gallego, Aina, Thomas Kurer, and Nikolas Bahati Scholl (2021). "Neither Left-Behind nor Superstar: Ordinary Winners of Digitalization at the Ballot Box". In: *Journal of Politics*. DOI: 10.1086/714920.

Humlum, Anders (2019). "Robot Adoption and Labor Market Dynamics". In: *Working paper*. URL: https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/5dcf78576d59eb44eac86f63/1573877848584/humlumJMP.pdf.

Im, Zhen Jie et al. (2019). "The "losers of automation": A reservoir of votes for the radical right?" In: *Research & Politics* 6.1. DOI: 10.1177/2053168018822395.

Innovation, Science and Economic Development Canada (2021). "Views of Canadians on Artificial Intelligence". In: *Report*. URL: https://www.ic.gc.ca/eic/site/112.nsf/eng/07663.html.

Jaimovich, Nir et al. (2020). "The Macroeconomics of Automation: Data, Theory, and Policy Analysis". In: *Economic Studies at Brookings*.

Kurer, Thomas and Aina Gallego (2019). "Distributional consequences of technological change: Worker-level evidence". In: *Research & Politics* 6.1. DOI: 10.1177/2053168018822142.