JANUARY 2020

Automation: A Guide for Policymakers

James Bessen
Boston University Technology & Policy Research Initiative

Maarten Goos
Utrecht University School of Economics

Anna Salomons
Utrecht University School of Economics

Wiljan van den Berge
CPB Netherlands Bureau for Economic Policy Analysis

B | Economic Studies
at BROOKINGS

Introduction

Advancing technologies are increasingly able to fully or partially automate job tasks. These technologies range from robotics to machine learning and other forms of artificial intelligence, and are being adopted across many sectors of the economy. Applications range from selecting job applicants for interviewing, picking orders in a warehouse, interpreting X-rays to diagnose disease, and automated customer service. These developments have raised concern that workers are being displaced by advancing automation technology. Indeed, over 18 recent studies predict job losses from new automation technologies, including some predictions of massive job losses (Winick 2018). A large literature on worker displacement[1] suggests that the effects of such developments could be dire: individual workers subject to plant closings and mass layoffs experience reduced employment probabilities and wage reductions, leading to long-term earnings losses, as well as reductions in consumption and worse health outcomes. Concerns about these effects of automation have led some commentators to call for policies to directly combat mass unemployment, such as a Universal Basic Income.[2]

But is this right? At a time when many firms are investing in automation, the unemployment rate is at historic lows. Low unemployment might seem hard to reconcile with apocalyptic predictions about mass unemployment. This paper reviews the evidence from recent studies and reports on a new paper we have written, "Automatic Reaction: What happens to workers at firms that automate" (Bessen et al. 2019). This paper is the first to take a look at what actually happens to those workers. We build on some of the findings in order to draw the implications for policy.

The evidence suggests that the apocalyptic predictions may have it wrong, at least for the next several decades. But that does not mean that automation is not highly disruptive. The challenge of automation in the near future may not be mass unemployment, but, instead, a greater level of worker transitions. Workers need to switch jobs, often learning new skills, changing occupations and industries, and moving to new locations. These transitions often involve temporary unemployment spells and a loss of income. Automation thus places a burden on workers even if, in the end, they do not permanently lose employment. Moreover, inefficient transitions can slow the productivity-enhancing promise of new technology.

Our analysis suggests that automation does indeed pose significant challenges for policymakers. But it is critical to respond to the real challenge, the challenge that workers are actually experiencing. This review of the evidence may be helpful toward that end.

[1] Starting with Ruhm 1991 and Jacobson et al. 1993.

[2] For example, Andrew Yang, see https://www.yang2020.com/what-is-freedom-dividend-faq/

Background

Much discussion of automation lacks a clear understanding of what automation is and how it works. This section reviews basic economic concepts regarding automation and recent empirical findings on the effects of new technologies on jobs.

Technology is about more than automation

First, it is important to note that technology does more than automate, by which we mean replace human work with work performed by machines. Technology creates new products that can disrupt or replace entire industries, as the automobile industry replaced the industry for horse-drawn carriages. Online advertising is certainly having a disruptive effect on traditional advertising media.

Also, technology can enhance human capabilities rather than simply displace them. Surgeons are able to perform better with robotic surgery tools; inventory management systems let retail buyers better predict demand for products; logistics systems let shipping managers better plan delivery of shipments. A recent survey of artificial intelligence startups asked what were the main capabilities this new technology provided to the startups' customers (Bessen et al. 2018). A significant number agreed that their products reduced labor costs and automated routine tasks. But many more reported that their products enhanced customer capabilities, including by making better predictions and decisions, by managing and understanding data better, and to create new capabilities to improve services or to provide new products.

These new capabilities are not necessarily socially benign. The disruption of industries might be of concern to policymakers. Also, even technologies that enhance capabilities may tend to increase economic inequality or to increase the market dominance of large firms, reducing competition (see Bessen 2018). There is no particular reason to assume that these effects pose any less of a policy challenge than does automation. Nevertheless, the focus here is on automation.

Just because something can be automated, doesn't mean it will be

Many of the predictions of job losses from automation cited above seem to assume that if a job is "automatable," then it will be automated. For example, Frey and Osborne (2017) began their analysis by having a panel of machine learning experts evaluate some 70 occupations to decide which were "fully automatable" in 2013. They used these judgments to project the total range of automatable jobs. But it is not clear that their basic assessments are accurate. They decided, for instance, that the jobs of accountant, loan officer, and judicial law clerk were fully automatable in 2013, given sufficient data. But, of course, these jobs were not actually fully automated then nor are they today.

The reason is that the technology has to be adopted and implemented before it automates anything. Economic and other considerations can prevent firms from adopting a new technology, for instance, if the technology costs more than the incremental profits it creates. This may be particularly true for early-stage versions of a technology that may perform rather poorly. Manyika et al. (2017) make some guesstimates about adoption rates, finding lower rates of displacement than Frey and Osborne. However, few other studies attempt such an exercise.

Furthermore, even if an automation technology is adopted, the impact on employment is not necessarily negative, as we develop below.

Tasks are automated, not jobs

A key distinction is that any particular automation technology tends to automate a particular task, not all the tasks in a job. Only rarely are entire jobs eliminated as the result of automation. Consider what happened to the 271 detailed occupations used in the 1950 Census by 2010 (Bessen 2016). Many occupations were eliminated for a variety of reasons, but few of these were automated away. In many cases, demand for the occupational services declined (e.g., boardinghouse keepers); in some cases, demand declined because of technological obsolescence (e.g., telegraph operators). This, however, is not the same as automation. In only one case—elevator operators—can the decline and disappearance of an occupation be largely attributed to automation. Nevertheless, this 60-year period witnessed extensive automation, but it was almost entirely *partial* automation. Any occupation involves many different tasks. While technologies may automate many tasks, it appears that technologies are not typically able to automate *all* the tasks needed to perform an occupation.

For example, economists have recognized that computers are particularly good at automating routine tasks.[3] Computers are often used to automate tasks that are repetitive and follow explicit rules. And so it is with occupations such as accountants and loan officers, jobs that have been partially automated since the 1950s. General Electric introduced a system to do payroll accounting in 1954; Fair and Isaac introduced an automated system to produce credit scores (the famous FICO score) to assist loan officers in 1956. And since then, many more technologies have automated additional tasks in these occupations. Yet despite the conclusion of Frey and Osborne's machine learning experts, the companies using the latest artificial intelligence technology to provide tools for accountants or loan officers do not claim to automate all of the tasks required for these jobs. The jobs are just too complex.

For example, while computerized systems work for issuing consumer credit cards where standardized data is available on consumer credit history, income, etc., they might not work so easily for other kinds of lending such as complicated business loans where the evaluation of credit risk depended on an ongoing assessment of business prospects. Simple standardized data about the firm is not enough. Instead, a loan officer needs to understand the risks that the business faces both at the time of the application and afterward, when the loan officer needs to monitor the business performance. For example, successful lenders

[3] See Bresnahan (1999), Autor, Levy, and Murnane (2003), Autor and Acemoglu (2011) and Goos, Manning, and Solomons (2014).

need to know when to call in a loan if a business appears to be failing. In order to make these assessments, loan officers gather a wide range of non-standardized information and apply judgment informed by experience and a wealth of background information. They may look at the business production facilities and human resource practices; they may talk to customers of the business to gauge their satisfaction and future demand; they will look at market projections and factor in overall economic forecasts; they will form opinions based on experience about the trustworthiness of the principals and the effectiveness of organizational controls. While automated tools might assist loan officers in some of these tasks, many of these tasks do not seem likely to be automated anytime soon.

Similarly, accountants and auditors perform many tasks that have not been automated and are not likely to be automated soon. They must assess the firm's assets, appraising their value and estimating risk and materiality (whether the variation in asset value significantly affects the firm's value) under a variety of scenarios; they must evaluate organizational controls that are used to limit misappropriation of funds and fraud; they must check reported transactions to verify that they have been stated accurately; they must interpret continually changing accounting regulations and tax laws, applying difficult judgments to the treatment of revenue, capitalized expenses, and more; they must report on the overall representation of the financial statements; they need to advise managers on how to improve accounting controls and financial performance. Finally, accountants and auditors still have to deal with paper. Despite claims of a paperless office in the 1980s, the use of paper in offices has quadrupled since then, only recently leveling off.[4]

Thus, the idea that accountancy and loan management jobs are fully automatable today seems unlikely. Nevertheless, it is certainly true that a growing number of tasks performed by these occupations are automatable with new technology. After all, that has been true for over 60 years. Arntz et al. (2016) attempt to make predictions of job losses based on automatability of tasks. They conclude that "cutting-edge digital technologies have little effect on aggregate employment."

In summary, to the extent that technology automates work, jobs tend to be only partially automated. The typical human worker in almost any occupation performs a wide range of tasks and even today's most highly sophisticated technologies cannot perform all of them. But even if jobs are only partially automated, jobs can be lost. However, as we explore next, this does not necessarily happen.

When industries automate, their employment can grow

Many people assume that if technology automates even part of a job, then employment will fall. One might come to that conclusion after looking at manufacturing employment. For example, there were 400,000 cotton textile workers in 1940 in the US;

[4] Christopher Mims, "Why the Paperless Office is Finally on Its Way," Wall Street Journal, September 18, 2016.

there are fewer than 20,000 today (see Figure 1). While trade is responsible for some of these losses in the last decade or so, automation accounts for most of the drop.[5]

Figure 1. US Textile Production Workers, cotton and synthetics

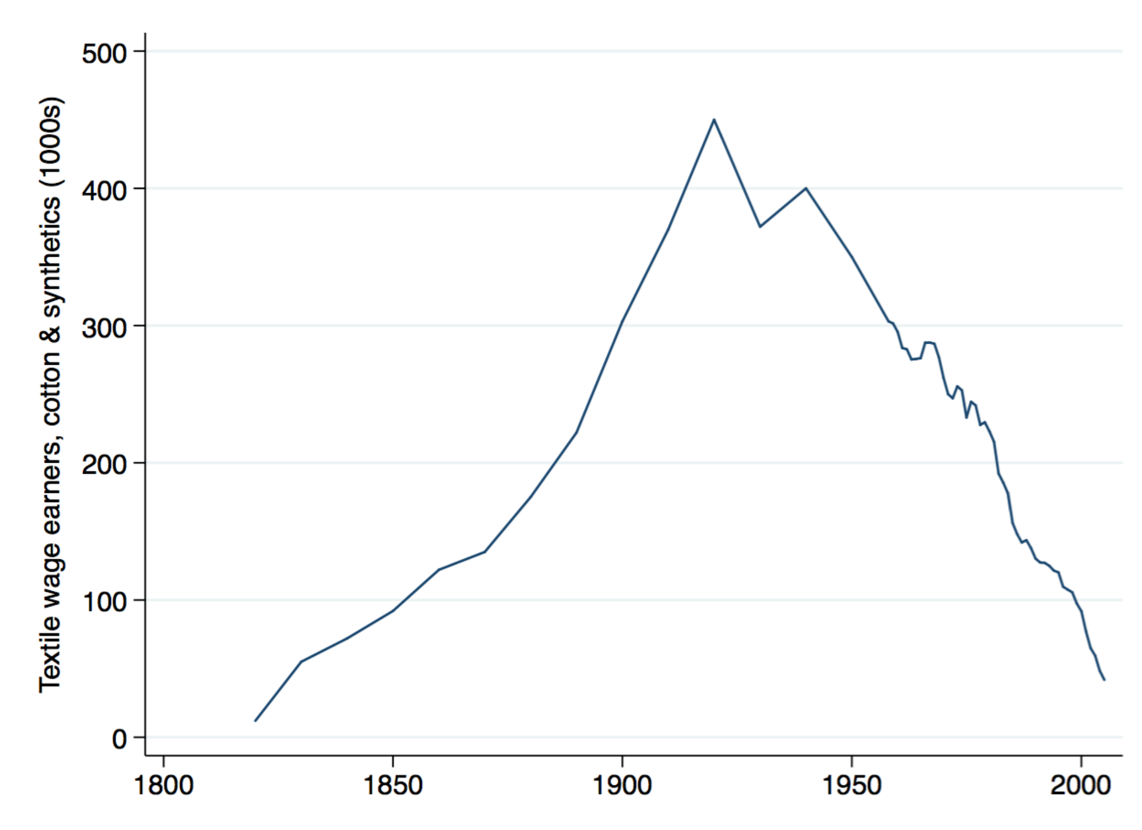

But it is a mistake to jump to the conclusion that automation of tasks leads to lower employment. In fact, for over 100 years prior to the 1940s, textile industry employment grew alongside relatively rapid automation (see Bessen 2019). Indeed, we didn't get to 400,000 textile workers without this growth. How could that happen? As expected, automation reduced the amount of labor required to produce a yard of cloth. But in a competitive market, this led to lower prices. And when there was much pent up demand—people had very few clothes or other textile items in the early 1800s—lower prices led to greater demand. So many more yards of cloth were demanded that total employment increased even though the labor required per yard fell. Automation induced a highly elastic demand response.

But by the mid-twentieth century, textile demand in the US was satiated. People had closets full of clothing, cloth upholstery, draperies, etc. Automation has continued to

[5] Until the last few decades, the balance of trade in cotton and synthetic textiles was not always negative and was relatively small.

increase output per worker at about the same rate, but now the demand response is inelastic. Demand no longer offsets the labor-saving effect of automation and so employment falls.

Additionally, this pattern is seen not just in textiles, but also in the primary steel and automotive industries, industries that also experienced sustained rapid productivity growth. These industries also experienced an "inverted U" in employment, growing for many decades at first and then levelling off or declining. Moreover, there is some reason to believe that this inverted U pattern might apply more generally to other industries (Bessen 2019).

Similar patterns can be observed with recent technologies. The automated teller machine took over cash handling tasks from bank tellers, but the number of bank tellers actually increased (Bessen 2016); the barcode scanner reduced the checkout time of cashiers in grocery stores by 18-19%, but the number of cashiers grew; electronic discovery reduced the time needed to search for documents in litigation, but the number of paralegals grew.

Thus, partial automation does not necessarily lead to job losses in the affected industries. At different times, the same industry can experience job growth or job losses with automation. At any one time, employment will grow in some industries and shrink in others. The nature of demand is critical to understand the employment response. Today, automation is beginning to affect a whole new set of industries in the service, finance, wholesale and retail sectors. It is an empirical matter to determine whether these industries have pent up demand—they have seen relatively little automation to date—or an inelastic demand response.

When industries cut jobs, unemployment doesn't necessarily go up

Even if automation reduces employment in some industries, that does not mean that overall employment decreases. The macro economy responds in ways that may offset job losses in some industries, creating job growth elsewhere. There are several ways this can happen. First, employment can grow in upstream or downstream industries in response to automation. For example, automation of textiles affects employment in cotton production and also in the production of clothing.

Second, there are income effects. Even if textile employment is reduced, national income increases because more goods are produced with the same resources. This income flows to owners and to consumers as well as to workers. Rising income increases demand in other industries aside from the demand effect in the automating industry. Autor and Salomons (2018) study the aggregate impact of productivity growth. They find that productivity growth often does decrease employment in the affected industry, but that the downstream and income effects offset this, leading to a rise in employment.

When tasks are automated, new labor-using tasks also arise

Much of the discussion on automation implicitly assumes the set of tasks performed in the economy as fixed: task automation, however partial, then implies human labor is confined to an ever-narrowing set of activities. However, empirical evidence supports the emergence of new activities (Lin 2011), and the importance of new labor-using tasks arising has been integrated in canonical models of automation (Acemoglu and Restrepo 2018b). These countervailing effects are only beginning to be studied empirically (Autor and Salomons 2019), but are important for thinking about how task automation shapes human work.

Recent evidence on employment effects

The above analysis argues that automation does not necessarily decrease employment at the affected firms, affected industries, or in the macroeconomy. But the actual impact of automation on employment is an empirical matter. Eleven recent studies have looked at the effects of potentially labor-saving technology on employment (see Table 1), both at firm and macro levels. Do these studies support the view of automation causing rising unemployment on a large scale?

Table 1. Automation Studies

Technology / Study	Country/Region	Employment effect
IT / AI		
Akerman, Gaarder, and Mogstad (2015)	NO	+ skilled workers
Bessen and Righi (2019)	US	+ services, FIRE, trade; - manufacturing
Gaggl and Wright (2014)	UK	+ trade and finance
Mann and Püttmann (2017)	US	+ services, - manufacturing
Automation		
Cirerra and Sabetti (2019)	53	0
Domini et al. (2019)	FR	+
Industrial Robots		
Acemoglu and Restrepo (2018a)	US	-
Dauth et al. (2017)	DE	0
Dixon, Hong, Wu (2019)	CA	+

Graetz and Michaels (2017)	DE	0
Koch, Manuylov, Smolka (2019)	SP	+ on adopting firm, - on non-adopting

There are two versions of the mass unemployment hypothesis one can consider. The strong form, corresponding to many of the predictions, holds that unemployment causes absolute employment losses in the affected industries or firms. The weak form holds only that employment grows in automating firms and industries more slowly than it would have grown without the automation.

Figure 2 shows evidence from our study (Bessen et al 2019) that rejects the strong form of the mass unemployment hypothesis. A similar result is found in many of the other micro-level papers. Automating firms tend to grow faster than other firms, both before and after the automation event. This finding cannot be easily reconciled with the view that automation causes absolute declines in employment.

Figure 2. Employment at firms making major automation investments and not, Netherlands (Bessen et al. 2019)

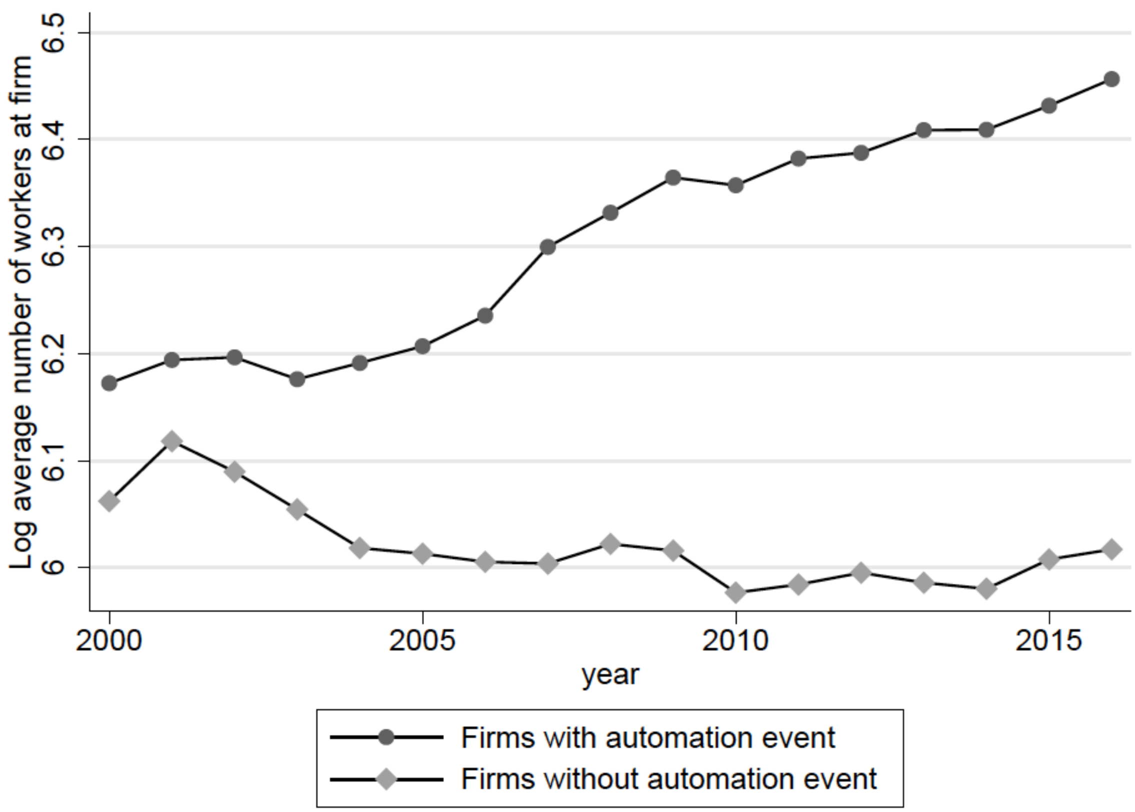

But even though employment rises faster at automating firms on average, it is entirely possible that it does not rise as fast as it would have otherwise. In this weak form of the mass unemployment hypothesis, automation slows employment growth, making the economy anemic.

The challenge in testing this hypothesis is that it is often difficult to specify what employment growth would have been without automation, the counterfactual "otherwise." The papers in Table 1 use a variety of econometric techniques to establish a counterfactual control, some more convincing than others. Nevertheless, the overall picture of the reported employment effects does not support the weak mass unemployment hypothesis. Only one of the twelve papers finds a substantial negative effects across all industries. Some of the papers find effects that differ by industry sector, with manufacturing tending to a more negative response.

Is this time different?

Even though there isn't strong evidence that automation is causing mass unemployment, perhaps things are about to change. Some people, such as Martin Ford (2015), argue that "this time is different." Perhaps the rate of technological change is much more rapid with new artificial intelligence technologies.

This is possible, but the role of demand in determining the employment outcome of automation provides a reason to be cautious about such predictions. First, if demand for the products being automated is highly elastic, then more rapid technical change will bring *faster* employment growth, not unemployment. This growth, requiring a reallocation of labor, might very well be disruptive, but not because of mass unemployment. Moreover, historically the elasticity of demand changes slowly (Bessen 2019), so the effects seen in Table 1 are likely to persist for at least a couple decades.

More generally, in the past, many people, have made predictions of imminent mass unemployment arising from automation. But these predictions have almost all failed to appreciate the depth of human demand. For example, in 1930, Keynes (1930), anticipating continued productivity growth, predicted that in 100 years his grandchildren would enjoy a fifteen-hour workweek. Now that we are close to that 100-year mark, the average workweek for OECD nations is 34 hours. Yet Keynes was right about productivity growth. In the US, the 1930 level of mean GDP per capita could be realized in 15 hours on average by 1977. What Keynes did not grasp was the depth of human wants and desires, that is, the depth of consumer demand. The reason we don't work 15-hour weeks is that we choose to demand more goods and services that technology has made cheaper and better.[6] A similar underestimation of demand lies behind many other failed predictions of automation-induced mass unemployment.

In the more distant future, of course, automation might cause mass unemployment or a substantially shorter work week. Perhaps future technologies will be able to fully automate more jobs; perhaps demand will become satiated in most markets; perhaps new industries and new occupations will no longer be created in the future. But such

[6] Also, many new products have enhanced our leisure time, including automating tasks such as washing clothes and dishes.

possibilities are a long way off. In the meantime, however, automation is nevertheless posing a considerable challenge to society. That challenge is not mass unemployment, but worker adjustments.

Impact on workers

Automation in the Netherlands

In order to measure the impact of automation on workers, we conducted a study (Bessen et al. 2019) using a unique dataset that captures annual firm spending on automation services. These data have been collected by the statistical authorities of the Netherlands since 2000 and are linked to administrative data on firms and on individual workers. There are few other sources of data on all types of automation at the firm level and few linked to the extensive firm and worker data needed to explore impacts. These data cover about 5 million workers annually from 2000 – 2016 at some 36,490 firms in all major private, non-financial sectors.

To analyze the impact of automation on workers, the study devises an innovative approach: we look at what happens when firms make major investments in automation. It turns out that much investment in automation occurs in discrete episodes of heavy investment, which we call “spikes.” These events allow us to compare the outcomes of workers at automating firms to a comparable “control” group of workers at similar firms that automate at a later time. We refer readers to the study paper for details of how we identify these spikes, why and how they occur, and how we construct our analysis to measure outcome effects. Here we summarize the main findings for the purpose of deriving policy implications.

The burden of automation

Unemployment and wages

We study the impact of automation on incumbent workers, those workers who were employed at their firm for three or more years before the automation event. Over the five years following automation, these workers lose, on average, about 11% of one year’s earnings or, in absolute terms, 3,800 euros.
These losses could arise from lower wages or from spells of non-employment. In fact, the daily wage rate does not change for these workers. These workers do not appear to experience reduced wage rates either if they stay at their firm or if they move. The lost earnings come from spells of non-employment attributed to automation. These total 18 days per worker on average (for both leavers and stayers) over five years.

Relatedly, we see an increase in the share of incumbent workers who leave their firms attributable to automation, although we do not know whether these workers were laid off or left of their own choice. During the year of the automation event, about 2% more of the incumbent workers leave by comparison to the control group. Over five years, the cumulative separations is less than 13%. The separations following automation appear to occur as a trickle rather than as a mass layoff.

These separating workers are offset by new hires at the firm, although it appears that the new hires do not fully offset the separations during the first years after the automation. However, our focus here is on the impact on incumbent workers regardless whether they are replaced or not. And the evidence shows a significant, although not overwhelming, loss of income and days worked. Further, these losses are only partially offset by benefits from the Dutch social safety net. Overall, incumbent workers recoup some 13% of their losses from unemployment benefits, disability benefits, and welfare payments. This finding is comparable to that in other worker displacement events, where typically only a small part of the negative impact is compensated by social security (Hardoy and Schøne 2014).[7]

Where do workers go?

Workers impacted by an automation event are more likely to switch industries (two-digit industries) than workers in firms that automate later. Besides industry switches, we do not find economically sizable or statistically significant changes in terms of workers' average or median firm wage, firm size, or firm automation expenditure. This implies that automation-affected workers are not structurally moving to firms that pay different wages, are different sizes, or are differently automation-intense. Affected workers are also somewhat more likely to take early retirement and to enter self-employment.

Who is affected?

We find that automation affects a large number of workers throughout the economy. In each year of our sample, about 9% of the incumbent workers are at firms that experience an automation event. While not every worker at each automating firm is affected, the number of workers at those firms is large.

And these firms are in all private non-financial sectors (finance is not included in our data). That is, automation is much broader than just industrial robots being used in manufacturing. While workers in manufacturing experience income losses from automation, so do workers in every other major sector except accommodation and food service. Also, while large firms tend to spend more on automation per worker than do smaller firms on average, many small firms do spend significantly.

We also look at worker characteristics. We do not find significant differences in outcomes by gender. Nor do we find significant difference by wage after controlling for worker age. That is, contrary to a common perception, it is not low wage workers who are primarily affected by automation. When we compare workers by wage to other workers in the same firms, we find, if anything, that higher wage workers experience larger relative income losses, although the difference is not statistically significant. Since our approach captures all automation technology, this could be consistent with Webb (2019), who uses

[7] In part, this is by law: unemployment benefits in the Netherlands have a replacement rate of 75 percent in the first two months of unemployment, which then decreases to 70 percent. Further, there is a maximum ceiling, such that workers with higher wages earn lower replacement rates than the 70 or 75 percent maximum.

patent data to show that while low-skilled workers are most exposed to robotics, other automation technologies such as software and artificial intelligence impact more on work performed by medium and high-skilled workers.

We do find that older workers are more severely affected. Workers 50 years of age and older experience larger income losses, largely due to longer periods of non-employment. Older workers appear to have a harder time finding new work.

Is automation like mass layoffs?

It is helpful to compare our findings to the literature on the effects of mass layoffs and plant closings. Economists have studied what happens to workers after these events that arise from bankruptcy, changing demand conditions, technological obsolescence and other reasons. Automation events tend to affect fewer workers by comparison to those events and workers leave more gradually after automation events. While the mass layoff literature studies layoffs of 30 percent of the workforce or more, the automation events only account for a 2 percent increase in incumbent employee separations during the first year followed by a continuing trickle of separations. Also, mass layoffs have been found to permanently reduce wage rates; we find no such effect for workers subject to major automation events.

We can do a back-of-the-envelope estimate of the total share of the workforce separating each year as the result of automation, at least as captured by our measure. Each year, about 9 percent of the workers in our sample are employed in firms that have an automation spike. Since the cumulative percent of incumbent workers who separate following an automation spike is less than 13 percent after 5 years, around 1 percent (.09 x .13) of treated incumbent workers leave their employers each year as a result of these spikes. In contrast, Abbring et al. (2002) find that about 4 percent of workers in both the US and the Netherlands are displaced each year because of their employers' adverse economic conditions.

Thus, while automation imposes significant losses on some workers, the impact seems somewhat smaller than the impact of mass layoffs and plant closures that occur from ongoing economic turbulence.

Implications for policy

A substantial portion of the workforce is now affected by automation each year and workers at automating firms experience non-negligible income losses, about 11% of a year's pay on average. This impact presents several policy challenges including distributional concerns about fairness and economic inequality. It also raises concerns about providing workers the skills, knowledge, and incentives to adjust to new jobs. Smoothing worker transitions is important both for worker welfare and also for society's ability to reallocate workers to best use productivity-enhancing new technology.

In a sense, the comparison to mass layoffs suggests that rising automation represents an increase in a challenge that society already faces, not a challenge of an entirely new sort. The ranks of workers leaving their employers following automation events adds to those workers who are adversely affected by mass layoffs and plant closings. Further, similar to lay-offs from Schumpeterian creative destruction, this job reallocation is a by-product of productivity growth, and workers would not be better off either in aggregate or on average without it.

But the evidence suggests that there is a role for governments to facilitate these transitions following automation episodes. The Dutch social safety net currently offsets about 13 percent of their losses. Most of these workers get new jobs after a spell of unemployment; some of these workers switch industries or become self-employed.

These findings imply that policymakers need to consider programs to ease worker transitions to new jobs, to new skills, new occupations, new industries and to new locations. These measures might include re-training, relocation assistance, and temporary income support. On the other hand, existing policies that discourage or hamper worker transitions, such as employee noncompete agreements, could be viewed as problematic.

In the absence of stronger evidence that automation is causing significant net losses in employment, our findings suggests that policymakers might do well to focus on helping workers transition to new jobs on an ongoing basis rather than to focus on a yet-to-be-seen mass unemployment from automation. The real burden of automation is already falling on some workers today, we can measure its impacts, and formulate policies to ameliorate them.

References

Abbring, J., van den Berg, G., Gautier, P., van Lomwel, A., van Ours, J., and Ruhm, C. (2002). Displaced Workers in The United States and The Netherlands. In Kuhn, P., editor, *Losing Work, Moving on: International Perspectives on Worker Displacement*. W.E. Upjohn Institute.

Acemoglu, Daron, and David Autor. "Skills, tasks and technologies: Implications for employment and earnings." Handbook of labor economics 4 (2011): 1043-1171.

Acemoglu, D. and Restrepo, P. (2018a). Robots and Jobs: Evidence from US Labor Markets. MIT.

Acemoglu, D. and Restrepo, P. (2018b). The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment. *American Economic Review*, 108(6):1488–1542.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. "The skill complementarity of broadband internet." The Quarterly Journal of Economics (2015), 1781–1824.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. "The risk of automation for jobs in OECD countries." (2016).

Autor, David H., Frank Levy and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration." Quarterly Journal of Economics, 116(4), 2003.

Autor, David, and Anna Salomons. "Is Automation Labor Share-Displacing? Productivity Growth, Employment, and the Labor Share." Brookings Papers on Economic Activity (2018): 1-64.

Autor, David, and Anna Salomons. "New Frontiers: The Evolving Content and Geography of New Work in the 20th Century". Mimeo (2019).

Bessen, James E. "How computer automation affects occupations: Technology, jobs, and skills." Boston Univ. school of law, law and economics research paper 15-49 (2016).

Bessen, James (2018), "The Policy Challenge of AI," CPI Antitrust Chronicle, June 2018.

Bessen, James E. "Automation and jobs: When technology boosts employment." *Economic Policy* (2019 forthcoming).

Bessen, J. E., Goos, M., Salomons, A., and Van den Berge, W. (2019). Automatic reaction-what happens to workers at firms that automate? *Boston Univ. School of Law, Law and Economics Research Paper*.

Bessen, James E., Stephen Michael Impink, Robert Seamans, and Lydia Reichensperger. "The Business of AI Startups." Boston Univ. School of Law, Law and Economics Research Paper 18-28 (2018).

Bessen, James E., and Cesare Righi. "Shocking Technology: What Happens When Firms Make Large IT Investments?." Boston Univ. School of Law, Law and Economics Research Paper 19-6 (2019).

Bresnahan, Timothy. "Computerisation and Wage Dispersion: An Analytical Reinterpretation." Economic Journal 109.456 (1999): F390-415.

Cirera, X. and Sabetti, L. (2019). The effects of innovation on employment in developing countries: evidence from enterprise surveys. *Industrial and Corporate Change*, 28(1):161–176.

Dauth, W., Findeisen, S., Südekum, J., and Wößner, N. (2017). German Robots – The Impact of Industrial Robots on Workers. Technical report, Institute for Employment Research, Nuremberg, Germany.

Dixon, Jay, Bryan Hong, and Lynn Wu. "The Employment Consequences of Robots: Firm-Level Evidence." Available at SSRN 3422581 (2019).

Domini, Giacomo, Marco Grazzi, Daniele Moschella, and Tania Treibich. Threats and opportunities in the digital era: automation spikes and employment dynamics. No. 2019/22. LEM Working Paper Series, 2019.

Ford, Martin. Rise of the Robots: Technology and the Threat of a Jobless Future. Basic Books, 2015.

Frey, Carl Benedikt, and Michael A. Osborne. "The future of employment: How susceptible are jobs to computerisation?." Technological forecasting and social change 114 (2017): 254-280.

Gaggl, Paul, and Greg C. Wright. "A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive." Working Paper (2014).

Goos, Maarten, Alan Manning, and Anna Salomons. "Explaining job polarization: Routine-biased technological change and offshoring." The American Economic Review 104.8 (2014): 2509-2526.

Graetz, Georg, and Guy Michaels. "Robots at work." Review of Economics and Statistics 100, no. 5 (2018): 753-768.

Hardoy, I. and Schøne, P. (2014). Displacement and Household Adaptation: Insured by the Spouse or the State? *Journal of Population Economics*, 27(3):683–703.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings Losses of Displaced Workers. *The American Economic Review*, pages 685–709.

Keynes, John Maynard. "Economic possibilities for our grandchildren." In Essays in persuasion, pp. 321-332. Palgrave Macmillan, London, 2010.

Koch, M., Manuylov, I., and Smolka, M. (2019). Robots and firms.

Koning, P. and Lindeboom, M. (2015). The Rise and Fall of Disability Insurance Enrollment in the Netherlands. *Journal of Economic Perspectives*, 29(2):151–72.

Lin, Jeffrey. "Technological adaptation, cities, and new work." Review of Economics and Statistics 93, no. 2 (2011): 554-574.

Mann, Katja and Lukas Püttmann, "Benign Effects of Automation: New Evidence From Patent Texts" working paper, 2017.

Manyika, James, Susan Lund, Michael Chui, Jacques Bughin, Jonathan Woetzel, Parul Batra, Ryan Ko, and Saurabh Sanghvi. "Jobs lost, jobs gained: Workforce transitions in a time of automation." McKinsey Global Institute (2017).

Ruhm, C. J. (1991). Displacement Induced Joblessness. *The Review of Economics and Statistics*, 73(3):517–522.

Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. Working paper, Stanford University.

Winick, E. (2018). Every study we could find on what automation will do to jobs, in one chart. *Technology Review*. https://www. technologyreview. com/s/610005/every-study-we-could-findon-whatautomation-will-do-to-jobs-in-one-chart.