

DISCUSSION PAPER SERIES

IZA DP No. 12428

**Digitalization and the Future of Work:
Macroeconomic Consequences**

Melanie Arntz
Terry Gregory
Ulrich Zierahn

JUNE 2019

DISCUSSION PAPER SERIES

IZA DP No. 12428

Digitalization and the Future of Work: Macroeconomic Consequences

Melanie Arntz

ZEW and University of Heidelberg

Terry Gregory

IZA and ZEW

Ulrich Zierahn

ZEW and CESifo

JUNE 2019

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Digitalization and the Future of Work: Macroeconomic Consequences*

Computing power continues to grow at an enormous rate. Simultaneously, more and better data is increasingly available and Machine Learning methods have seen significant breakthroughs in the recent past. All this pushes further the boundary of what machines can do. Nowadays increasingly complex tasks are automatable at a precision which seemed infeasible only few years ago. The examples range from voice and image recognition, playing Go, to self-driving vehicles. Machines are able to perform more and more manual and also cognitive tasks that previously only humans could do. As a result of these developments, some argue that large shares of jobs are “at risk of automation”, spurring public fears of massive job-losses and technological unemployment. This chapter discusses how new digital technologies might affect the labor market in the near future. First, the chapter discusses estimates of automation potentials, showing that many estimates are severely upward biased because they ignore that workers in seemingly automatable occupations already take over hard-to-automate tasks. Secondly, it highlights that these numbers only refer to what theoretically could be automated and that this must not be equated with job-losses or employment effects – a mistake that is done often in the public debate. Thirdly, the chapter develops scenarios on how digitalization is likely to affect the German labor market in the next five years and derives implications for policy makers on how to shape the future of work. Germany is an interesting case to study, as it is a developed country at the technological frontier. In particular, the main challenge will not be the number, but the structure of jobs and the corresponding need for supply side adjustments to meet the shift in demand both within and between occupations and sectors.

JEL Classification: J23, J31, O33

Keywords: automation, digitalization, unemployment, inequality

Corresponding author:

Terry Gregory
IZA - Institute of Labor Economics
Schaumburg-Lippe-Str. 5-9
53113 Bonn
Germany
E-mail: tgregory@iza.org

* Prepared for the Handbook of Labor, Human Resources and Population Economics (edited by Klaus F. Zimmermann).

1 Introduction

The past decades have been characterized by a tremendous rise of computing power. Since 1945, computing power increased, on average, by 45 percent per year, implying a drastic decline of the costs of computational tasks (Nordhaus, 2007). These rapid improvements have been accompanied by computer-controlled automation of so-called routine tasks. Routine tasks are tasks which follow well-defined rules and can thus be automated based on rule-based algorithms, using rapidly improving computers. As a consequence, labor demand for routine tasks has generally declined. As routine tasks were wide-spread among many middle-skilled, medium-wage workers, such as bookkeepers, clerical assistants or production workers, this computerization has led to a polarization of the labor market in the recent past with declining shares of middle- and rising shares of both high- and low-wage workers (see Acemoglu and Autor 2011 for a review of the literature).

While computerization has replaced humans in many routine tasks, these tasks have in common that they need to be well-defined. However, people understand many tasks tacitly, without being able to clearly pin down the exact underlying rules, limiting the scope of what can be automated based on software algorithms (“Polany’s Paradox”, Autor 2015). More recently, these technological barriers are reduced by Machine Learning (ML) methods, in particular by Deep Learning. Such methods are based on the idea of training machines in performing tasks by providing them with suitable data, instead of developing algorithms of well-defined rules. The machines “learn” how to do the task by mimicking the observed behavior, which implies that there is no more need for explicitly understanding the precise rules underlying the observed patterns.

Deep Learning has been advocated already decades ago under different names. It has been discussed as cybernetics in the 1940s-1960s, as connectionism in the 1980s-1990s, and as Deep Learning since about 2006 (Goodfellow et al., 2016). However, it has increasingly been applied in real-world settings only recently (see Brynjolfsson et al. 2019 for examples). This is due to the fact that these methods require both high computing power and large amounts of training data, both of which became increasingly available during the last years. Moreover, data collection and availability has become ubiquitous e.g. via smartphones or the internet. Therefore, the range of problems that can be solved via ML has extended significantly. This now allows for the automation of cognitive tasks that previously only humans could do, sometimes even exceeding human precision (Brynjolfsson et al., 2019). These include also tasks typically requiring high-skilled workers, see e.g. Brynjolfsson and McAfee (2016) or Pratt (2015).

Against this background, there exists a debate on how many tasks or occupations might be automatable in the near future. A corresponding study which received widespread attention in the public debate is Frey and Osborne (2017), who claim that about half of the US workforce are “at risk of automation” in the next one to two decades. While some authors and consultancy agencies make similar claims (e.g. Bowles, 2014, Pajarinen and Rouvinen, 2014, PWC, 2018), other authors report much lower figures (Arntz et al., 2016, Arntz et al., 2017, Nedelkoska and Quintini, 2018, Dengler and Matthes, 2018, Pouliakas, 2018).

This debate is accompanied by widespread public fears of technological unemployment. Historically, such fears are not new. In fact, it has been claimed many times in history that technological change will lead to mass unemployment (see Mokyr et al., 2015 for a discussion) or will even herald the “End of Work”, as for example popularized by the eponymous book by Rifkin (1995). So far, these fears have not come true, raising the question of why there still are so many jobs (Autor, 2015).

In order to shed light on this question, this chapter discusses recent evidence on automation potentials and how this might translate into actual employment effects, see also Figure 1. The




Figure 1: From New Technologies to Their Employment Effects

chapter starts out by discussing how many and which kind of jobs might be automatable in the near future (Section 2). The focus lies on explaining the large differences in corresponding estimates. Secondly, Section 3 debates what this might imply for employment. The latter is important, as automation potentials only capture technological feasibility to automate jobs, which must not be equated with actual employment effects, since the diffusion of new technologies in the labor market is a slow and incomplete process (Section 3.1), workers adapt (Section 3.2) and potential job destruction effects might be compensated by job creation effects (Section 3.3). This distinction between automation potentials and actually resulting employment effects is often ignored in the public debate. As large automation potentials already existed in the past without resulting in mass unemployment, the chapter will also look back in time to see why this hasn't been the case and what one can learn from the past for the future of work. Thirdly, based on new and unique data on the use of digital technologies in the German economy, Section 4 presents first estimates of how automation via digitalization might affect the German labor market in the next few years. Section 5 concludes.

2 Automation Potentials

How many and which jobs are susceptible to future automation? In order to address this question, Frey and Osborne (2017) use the following approach: They ask experts in ML what machines are able to do and extrapolate their assessments to the U.S. workforce. At first, they subjectively hand-label 70 occupations from the O*NET database in a joint workshop with ML experts as either automatable or non-automatable. The O*NET database provides the task and job descriptions for each occupation. The question to be answered during the workshop was: “Can the tasks of this job be sufficiently specified, conditional on the availability of big data, to be performed by state of the art computer-controlled equipment” (Frey and Osborne, 2017, p. 263). Using this approach, the authors ultimately calculate the technical possibility of automating a job, which we refer to as automation potential. Note that the automation potential does not capture the probability that a job is actually automated, let alone the resulting employment effects of automation.


Frey and Osborne (2017) then rely on a selective list of variables regarding occupational tasks from the O*NET as well as the hand-labelled occupations from the workshop, to train an ML algorithm that classifies occupations as automatable or non-automatable. In other words, they estimate a statistical model of automation potentials using nine tasks indicators as explanatory variables. Finally, they use this model to extrapolate automation potentials for all 702 occupations that are included in the O*NET task data. The model returns an estimate of the automation potential. This number ranges between 0 and 100%. It is the likelihood that an occupation is technically automatable or, strictly speaking, it is an estimate of the probability that the experts would have classified a given occupation as automatable during the workshop. Frey and Osborne (2017) then define an occupation as automatable or as “at high risk”, if the ML algorithm returns at least an estimated automation potential of 70%. Finally, they combine this with occupational employment data to compute that 47% of workers in the U.S. are currently working in “high risk” or automatable occupations.

Several authors apply this approach to other countries by assigning the estimated automation potential for an occupation from Frey and Osborne (2017) to the country-specific occupational structure. Thereby, these studies assume that occupations are comparable across countries regarding their task structure. For example, Bowles (2014) finds that, on average, 54% of workers in the European Union are “at high risk”, with estimates ranging from 47% in Sweden to 62% in Romania. Pajarinen and Rouvinen (2014) argue that 36% of workers in Finland are automatable.

A key drawback of the approach by Frey and Osborne (2017) is that they focus on the occupational level, thus assuming all workers of the same occupation to conduct exactly the same tasks as described in the O*NET data. This is a very strong assumption as tasks do not only vary between workers of different occupations, but also vary substantially between workers of the same occupation (Autor and Handel, 2013). Moreover, the overwhelming majority of the decline in routine tasks in the context of computerization has been due to declining shares of routine tasks within occupations instead of declining shares of routine occupations (Spitz-Oener, 2006). Hence, to the extent that average occupational task structures do not sufficiently represent the task heterogeneity within occupations, especially regarding new, less automatable tasks, occupation-level approaches, such as Frey and Osborne (2017), are likely to overestimate automation potentials.

For this reason, Arntz et al. (2016, 2017) instead follow a different approach by focusing on what people actually do in their jobs rather than relying on occupational descriptions of jobs. For this, they use individual-level survey data provided by PIAAC (Programme for the International Assessment of Adult Competencies). Based on a statistical model that links the estimated automation potential by Frey and Osborne (2017) to the job-level characteristics of the workers in the PIAAC, they then show that only 9% of all U.S. workers are conducting automatable jobs (i.e. jobs with an estimated automation potential of at least 70%). Moreover, they calculate similar figures for other countries for which the survey is available and find that automation potentials vary between 6% in South Korea and 12% in Germany (see Figure 2).

Further differentiating the results by educational attainment, they find that low-skilled workers are particularly exposed to automation (Arntz et al., 2016). The share of workers with high automation potentials is highest for unskilled workers and strongly declines with educational attainment (see Figure 3). Similarly, low-income workers are exposed the most, whereas high-wage earners are least exposed to being potentially automatable. Hence, even though new automation technologies are increasingly capable to perform tasks of highly skilled and highly paid workers, it’s the low-skilled workers whose tasks are most exposed to being potentially automatable. This resembles the skill-biased technological change before the 1980s which favored higher skilled workers at the expense of lower skilled workers (Acemoglu and Autor 2011) rather



Source: Arntz et al. (2016)

Figure 2: Automation Potentials by Country

than the period from the mid-1980s until the 2000s during which computer-based technologies mostly substituted for middle-skilled routine-task-intensive workers.

Why do they find, on average, lower estimates compared to those of Frey and Osborne (2017)? The differences could stem either from the different level of analysis (occupation- vs. job-level), or from differences in data and methodology. In order to test the different explanations, Arntz et al. (2017) use occupation-level median task structures from the PIAAC to predict the automation potentials based on their estimated model. In this case, 38% of all U.S. workers fall in the “high-risk” group, implying that data and methodology can only explain a small part of the differences between the occupation-level and job-level approach (see Figure 4). Instead, much of the difference is explained by the huge variation of workers’ tasks within occupations. Apparently, workers in seemingly automatable occupations specialize in different types of hard-to-automate tasks such that occupational means do not represent the relevant task spectrum very well. As a result, occupation-level approaches, such as Frey and Osborne (2017), overestimate automation potentials.

A related study by Nedelkoska and Quintini (2018) also applies a job-level approach by using the PIAAC. In contrast to the Arntz et al. (2016, 2017) studies, they have access to 440 detailed ISCO occupations. Whereas Arntz et al. (2016, 2017) rely on the estimated automation potentials of all 702 occupations, they only use the expert evaluations of the 70 hand-labelled occupations from the workshop and link them to selective task variables in the PIAAC. Their methodology thus is closer to Frey and Osborne (2017). Nevertheless, using job-level information on the tasks conducted, they find 10% of U.S. workers to be in the “high-risk” group – as compared to the 47% by Frey and Osborne (2017). Pouliakas (2018) adopts a very similar approach, but instead uses another data set, the European Skills and Jobs Survey (ESJS), and finds that 14% of workers in the European Union work in automatable jobs. Hence, these studies confirm that differences between the results by Arntz et al. (2016, 2017) and Frey and Osborne




Figure 3: Automation Potentials by Education




Figure 4: Automation Potential Estimates

(2017) are not due to methodology or data, but are explained by the fact that occupation-level approaches severely overestimate automation potentials.

Other studies adopt different methodologies to estimate the share of automatable jobs. For example Dengler and Matthes (2018) hand-classify 8.000 different tasks from the German BERUFENET database, an occupation-level database similar to O*NET, as either automatable or non-automatable. They then calculate the share of automatable tasks for 4.000 different occupations and find that 15% of all workers in Germany are in the “high-risk” group, a finding that is quite comparable to the 12% for Germany by Arntz et al. (2016). A potential explanation for the similar finding is that they use such a disaggregated level of occupations that the corresponding occupational descriptions come close to describing the actual job-level. Another study by Arnold et al. (2016) conducts a survey among German employees, focusing on workers’ subjective beliefs about their jobs’ automation potential. According to their results, 13% of all respondents consider it likely or very likely that their job will be replaced by new technologies in the next 10 years.

Overall, there now exists widespread evidence that occupation-level studies overestimate automation potentials. Taking into account job-level variation of tasks within occupations, the share of automatable jobs drops to about 9% for the U.S. and to comparable figures in other countries. Nevertheless, the insights that can be drawn from such estimates remain limited, as the estimated automation potentials only capture whether a job – given its contemporaneous task structure – could theoretically be done by a machine or not. They remain silent about actual job losses or employment effects in the next two decades.

3 Automation and Employment

There are three main reasons why the previously discussed automation potentials must not be equated with actual or expected job losses or employment effects, see also Arntz et al. (2016) for a discussion:

1. Technological diffusion, i.e. the gap between technological potential and its actual implementation.
2. Worker flexibility, i.e. the ability of workers and jobs to adjust their tasks to new requirements.
3. Induced job creation, i.e. the job creation that technological change induces via several mechanisms.

Taking all three aspects together, the actual net employment effect of new technologies may surface with a time lag and may even be positive rather than negative. This section discusses all three aspects in more detail.

3.1 Technological Diffusion


In 1987, Solow famously wrote that “[you] can see the computer age everywhere but in the productivity statistics” (Solow, 1987). This quote became famous as the “Solow Paradox”. According to Brynjolfsson et al. (2019), we currently experience a comparable paradox with Artificial Intelligence (AI). AI systems rapidly advance and already surpass humans in selected tasks, but productivity slows down rather than rises. While different explanations can be put forward, they argue that large lags in AI implementation are the main contributor to the paradox. So far, AI adoption severely lags behind its technological capabilities. In particular,

Brynjolfsson et al. (2019) explain the slow diffusion of AI by arguing that AI is a General Purpose Technology (GPT). GPTs are technologies which are pervasive (i.e. they spread to most sectors), they improve over time, and they enhance the possibilities to further invent and produce new products and processes, such as for example electricity or information technology (Jovanovic and Rousseau, 2005).

GPT often require a long time to diffuse in the wider economy. For example, computers took 25 years to reach their long-run plateau of 5% of nonresidential equipment capital. About half of U.S. manufacturing plants remained non-electrified 30 years after the introduction of the polyphase system. These GPTs achieved widespread productivity gains and adoption only once sufficient complementary innovations and investments were made (Brynjolfsson et al., 2019). The rather low speed of diffusion is evident also for the latest technological advances. For instance, in the recent IAB-ZEW Labour Market 4.0 (LM4.0) firm survey, Arntz et al. (2019a) collected information on the level of technology underlying the capital stock that is used in German firms. In particular, they distinguish between manually controlled technologies that correspond to technologies that are either functioning mechanically or electrically, but are not IT supported, i.e. “1.0/2.0-technologies”, technologies that are supported by computers and software algorithms, i.e. “3.0-technologies”, and “4.0-technologies” that correspond to technologies that are IT-integrated, i.e. they allow for a direct and automated communication between different parts of the value chain such that workers only need to intervene in case of failures. In manufacturing, a production based on these highly automated, digital technologies is often referred to as “Industry 4.0”, echoing the fact that the underlying technological advances have been considered to constitute a new, fourth industrial revolution.

Figure 5 shows the technological structure of the capital stock used in German firms in 2016 at the time of the survey, the retrospective structure as of 2011, and the firms’ expected structure as of 2021, differentiated by production as well as office and communication equipment. For both types of the capital stock, the shares of capital based on 4.0-technologies roughly double in the ten year period. However, its share still remains small and, in fact, many firms are upgrading from 1.0/2.0 to 3.0 technologies rather than introducing the latest technologies. Thus, despite an ongoing diffusion of 4.0-technologies at a notable speed, the capital stock will continue to be dominated by older technologies in the near future. Or, in the words of Brynjolfsson et al. (2019, p. 10), “it takes time to build the stock of the new technology to a size sufficient enough to have an aggregate effect”.

Several reasons can be put forward to why the speed of diffusion may actually be lower than often expected. First of all, automation technologies will only be adopted if they can execute a particular task at lower costs than a worker. This is highlighted in the seminal framework on automation and jobs by Acemoglu and Restrepo (2018d). This framework highlights that what matters for the labor market is not how much could theoretically be automated, but how much of these technological capabilities are actually profitable to be adopted. Hence, the speed of diffusion also hinges on the costs of labor and thus on wage setting institutions such as minimum wages or the role of collective wage bargaining. While low labor costs at the lower end of the distribution may thus shield workers from being automated by machines to some extent, more expensive workers in the middle of the wage distribution may be more profitable to automate. At the same time, it is not clear when machines will actually have a comparative advantage to perform the more complex tasks at the upper end of the wage distribution. This will, of course, also depend on how much the wage level decreases in response to an automation-induced decline in labor demand, as the decline in wages improves worker’s employment prospects again. Acemoglu and Restrepo (2018b) for example model the effects of automation when workers have different skills. For this reason, technological capabilities do not necessarily translate into technological obsolescence of human labor (Acemoglu and Restrepo,



Source: Arntz et al. (2019a)

Figure 5: Technological endowment of German firms

2018d).

The costs of implementing new technologies are not limited to the acquisition of these technologies. Instead, additional investments often are necessary to fully utilize the new technologies and make them profitable. This is particularly true for major innovations such as GPT. Complementary investments comprise, for example, necessary organizational restructuring or the acquisition of the right skills via further training and new hires (Brynjolfsson et al., 2019). Firms go through a process of organizational redesign and substantially change their service and product mix to raise service quality and gain efficiency (Bresnahan et al., 2002). In fact, the shortage of qualified personnel that is able to handle new technologies may slow down its implementation, as the introduction of new technologies likely requires the availability of complementary skills (Acemoglu, 1998). In line with this, firms from the German LM4.0 survey, consider lack of qualified personnel to be a major risk for the implementation of new technologies. Moreover, the 65% of German firms which did not invest in 4.0-technologies between 2011 and 2016 particularly stress such risks and downplay the chances compared to firms which already use these technologies (Arntz et al., 2018). The barriers to the implementation of new technologies thus seem to be severe for a large share of firms, notably the smaller, less knowledge-intensive firms.

Apart from business-related reasons, there may also be ethical or legal obstacles that slow down or limit the speed of technological adoption. As a prominent example that has been discussed by Thierer and Hagemann (2015) and Bonnefon et al. (2016), the autonomous car bears new legal challenges regarding, for instance, the liability in case of an accident. Moreover, ethical questions emerge whenever an autonomous car cannot prevent an accident and an algorithm has to decide, for example, between crashing into a car or a truck. While some of these obstacles may be resolved at some point, they clearly slow down the pace with which technologies are introduced.

A final aspect that should be considered is that society may have strong preferences for the provision of certain tasks and services by humans as opposed to machines. As an example,

nursing or caring for the elderly may remain labor-intensive sectors, even if service robots increasingly complement these professions in the future. Hence, “*some human services will probably continue to command a premium compared to robotically produced one*” (Pratt, 2015, p. 58), meaning that there is a societal value attached to humans performing certain tasks that tends to preserve their comparative advantage.

3.2 Flexibility of Workers

Jobs are bundles of tasks, not all of which can be automated. Just because a certain fraction of a job’s tasks can be automated, the job need not be automated as a whole (see e.g. Brynjolfsson and Mitchell, 2017, Autor, 2015). For example, Arntz et al. (2017) argue that workers in seemingly automatable occupations apparently specialize in non-automatable niches within their profession. Most jobs are unlikely to be sufficiently well defined to be fully substitutable by machines. In line with this, Pratt (2015, p. 52) states for the advances in robotics that “*specialized robots will improve at performing well-defined tasks, but in the real world, there are far more problems yet to be solved than ways presently known to solve them*”.

Therefore, when firms introduce new production technologies, the initial impact of those machines on employment depends on whether workers are able to adjust to the new demands. In particular, the new technologies typically substitute for certain tasks and complement others. Whether automation technologies replace workers in a given job thus hinges on workers’ ability to exchange tasks that are replaceable by machines for new tasks that complement machines. For example, Automated Teller Machines (ATM) directly took over the tasks that bank tellers previously did in banks. Nevertheless, with the rapid increase in the number of ATMs, the number of bank tellers raised instead of declined. While ATMs replaced bank tellers in some tasks, bank tellers became more valuable in their remaining tasks, such as handling small business customers (Bessen, 2015).

Several studies suggest that this adjustment mechanism may actually be quite effective. For instance, although there has been a decline in jobs with predominantly routine and automatable tasks, the reduction of routine and automatable tasks in the economy mainly takes place by adjusting the set of tasks within occupations (e.g. Autor et al., 2003, Spitz-Oener, 2006). Workers seem to shift worktime from routine and automatable tasks to tasks that complement machines. The computerization for example has been associated with a strong decline in routine tasks. Spitz-Oener (2006) finds that less than 1% of the decline in cognitive routine tasks between 1979 and 1999 in Germany occurred between occupations, i.e. due to declining shares of cognitive-routine intensive occupation. Instead, almost all of the decline took place within occupations – i.e. workers in cognitive-routine intense occupations switching from cognitive-routine tasks to other tasks. More broadly, she finds the vast majority of task changes to take place within rather than between occupations.

Typically, the adoption of new technologies comes along with a new division of labor where workers increasingly perform tasks that complement machines (Autor, 2015), some of which may actually be newly created tasks (Acemoglu and Restrepo, 2018d). The tasks done by long-established occupations such as secretaries, for instance, clearly changed dramatically across time as skill demands changed with the introduction of new machines (e.g. typewriter, personal computer, workflow systems).

The adoption of new technologies is likely to differently affect workers depending on their abilities. Janssen and Mohrenweiser (2018), for example, investigate the introduction of the new computer-based control system (Computerized Numerical Control, CNC) in the field of cutting machine operation into the German apprenticeship regulation. The subsequent adoption of this technology likely was harmful for workers who graduated just before the change of the

curriculum. However, Janssen and Mohrenweiser (2018) find that only those workers who were forced to switch their occupations experienced negative labor market effects, suggesting that those who remained employed in the occupation learned to handle the technology on the job, shielding themselves from potential negative consequences. Cortes (2016) found that workers in routine occupations, who were exposed to computerization, experienced a wage increase of 14-16% over 10 years if they were able to switch to higher-paid cognitive jobs compared to those who stayed. These, however, potentially were the high-ability workers. Hence, workers' fate in phases of technological turmoil depends on workers' ability to learn the skills required in their new work environment, or on their ability to upgrade their occupation.

Overall, new technologies are unlikely to fully automate workplaces or occupations on a large scale, but rather change workplaces and the tasks involved in certain occupations. As long as workers are able to adjust to these new task demands, machines need not crowd out workers. However, if the tasks that complement machines become increasingly complex and demanding, the employment prospects for workers lacking certain skills may deteriorate.

3.3 Compensatory Mechanisms

Despite the just described adjustment of tasks at the level of workers and individual workplaces, the introduction of automation technologies to some extent replaces workers who were previously employed to perform the automated tasks. Whether this leads to an overall increase or decline of employment, is ambiguous, as several compensating mechanisms counteract the initial displacement effect. Acemoglu and Restrepo (2018d,c) develop a framework to analyze under which conditions the displacement effect of automation exceeds its compensating mechanisms. In particular, automation induces the following effects:

- *Productivity effect.* This effect captures the fact that technological innovations make firms more productive, reducing costs and prices which raises demand and production. In addition, automation may raise quality or enable new types of products or services, increasing demand and production. Moreover, the economy expands, raising the demand for labor also in sectors that do not adopt new technologies due to a multiplier effect.
- *Reinstatement effect.* This effect evolves either because these new tasks are complementary to the new technologies or because the displacement effect increases the amount of labor that is available to perform new, more productive tasks. More workers are required to perform the new tasks, raising demand for labor.

The net effect of automation on employment is ambiguous and ultimately remains an empirical question. There exists a huge empirical literature that analyzes the aggregate employment effects of new technologies and innovations, as for example surveyed by Feldmann (2013), Pianta (2009) or Vivarelli (2007). However, this chapter focuses on technologies that aim at substituting for workers via automation. This chapter solely discusses the empirical literature that explicitly addresses automation technologies, and particularly focuses on more recent technologies.

Studies on the employment effect of automation technologies can be roughly categorized into firm-, sector- and regional-level studies. Firm-level studies analyze how the adoption of automation technologies affects employment in the firms that invest in them. Cortes and Salvatori (2018), for example, do not find a decline of routine occupations in firms which invested in new technologies, contrary to what one would expect given that computerization potentially substitutes for routine tasks. Instead, they find that most of the decline in routine occupations is linked to declining shares of firms with initially larger shares of routine occupations. Similarly, in ongoing work Arntz et al. (2019a) show that firms' technology investments did not reduce


their net employment, because displacement effects are offset by technology-induced firm-growth in Germany. While the net effects at the firm level thus seem to be small, Bessen et al. (2019) show that automation in firms raises the separation rates of their workers, reducing their days in employment and wage income over the next five years. While firm-level studies are informative about the processes taking place within the firms, they remain rather silent about adjustment processes that occur between firms. In particular, if firms automate, they potentially become more competitive and crowd out firms that do not automate. The firm-level results therefore cannot be transferred to aggregate employment effects, as potential positive employment effects in automating firms could be offset by employment losses in competing firms.

Sector-level studies take into account this reallocation of workers between less and more innovative firms. A recent study by Graetz and Michaels (2018) for 17 OECD countries for example shows that the additional use of robots between 1993 and 2007 raised both labor productivity and value added at the sectoral level by about 0.36 and 0.37 percentage points, respectively, as suggested by the *productivity effect*. At the same time, they find no significant effects on total hours worked, although they report negative effects for low-skilled workers. Similar to firm-level studies, sector-level studies are only suggestive for the aggregate employment effects as they typically do not take into account the technology-induced reallocation between innovative- and non-innovative sectors.

Other studies rely on regions as small economies to study economy-wide effects of technological change. Dauth et al. (2017) find net neutral effects of robots in German local labor markets between 1994 and 2014. This is accompanied by a loss of about 2.12 jobs in manufacturing per additional robot, which is fully compensated by rising service employment. Hence, local labor markets with a higher exposure to robots did not experience net employment losses. Acemoglu and Restrepo (2017), to the contrary, document negative overall effects of robots in US local labor markets between 1993 and 2007. According to them, an additional robot per 1,000 workers reduces the employment-to-population ratio by 0.2 percentage points and thus has only limited effects on the US labor market.

Hence, *productivity* and *reinstatement effects* of robots apparently are strong enough to compensate their *displacement effects* in Germany, but are somewhat weaker in the US. There exist several potential reasons for this divergence. In Germany, labor protection legislation is stricter than in the US, meaning that it is more costly for German firms to lay off workers, raising firms' incentives to train workers to take over new tasks, rather than laying them off. In addition, the strong vocational education in Germany likely ensures that workers are higher skilled and better able to take over new tasks, compared to US workers. Finally, the higher formal education of the German workers exposed to robots in manufacturing likely implies that they might rank higher in the German wage distribution than comparable workers in the US wage distribution. In this case, the productivity gains from automation via robots are higher in Germany than in the US, such that the compensating mechanisms might be stronger in Germany compared to the US.

Acemoglu and Restrepo (2017) highlight that employment effects of robots seem to strongly differ from that of computerization more broadly. Gregory et al. (2018) instead study the employment effects of computerization. In contrast to the other two studies, they adopt a structural approach, by which they are able to explicitly disentangle the job destruction effects of computerization from the compensating mechanisms. They find that computerization indeed had strong displacement effects, reducing employment by 1.64 million jobs between 1999 and 2010 in the European Union. However, computerization created more additional jobs via induced productivity effects, resulting in a net employment increase of 1.79 million jobs (see Figure 6). Other results by Autor and Dorn (2013) indicate no net negative employment effects of computerization in the US. Differences in the effects of computerization in the European



Source: Gregory et al. (2018)

Figure 6: Employment Effects of Computerization in the European Union, 1999-2010

Union and in the US might result from the same reasons, as with robots above.

Evidence from recent automation phases thus suggests that there have been no negative employment effect of computerization in the US and even positive effects in the European Union. Robots instead, a technology that is more focused on replacing human tasks, do not, on net, destroy jobs in Germany, while they do reduce employment in the US, although to a limited extend only. According to this, previous automation technologies indeed did displace many workers, but had no or only limited negative employment effects due to large compensating mechanisms. Obviously, past phases of technological change so far did not lead to mass unemployment due to countervailing effects (Autor, 2015). It is therefore misleading to simply focus on automation potentials when one aims to understand how automation technologies affect the labor market.

4 Scenarios for Employment Effects

The previous sections highlight that automation potentials are not informative about the impact of automation on the labor market, as they ignore slow and incomplete technology adoption, worker level adjustment, and job creation effects. Several recent studies overcome these problems by directly studying the link between past automation and its employment effects. However, these studies focus on past automation technologies and do not capture the more recent technological innovations in the field of artificial intelligence. To overcome this shortcoming, Arntz et al. (2018) exploit a more recent period of technological upgrading in the German labor market (2011-2016) in order to simulate scenarios for the next few years (2016-2021) regarding the effects so called 4.0 technologies on employment and its various compensation mechanisms. In contrast to previous studies, they are able to take into account cutting-edge automation technologies. This section, first outlines the methodology, before presenting the results for the baseline and alternative scenarios and deriving implications.

4.1 Methodology

The basic idea of the methodology by Arntz et al. (2018) is to first estimate how technology investments have affected employment in the investing firms and the wider economy and to then study the likely consequences of more investments into cutting-edge technologies in the future. To do this, they first develop a task-based framework which captures mechanisms that are similar to mechanisms in the framework by Acemoglu and Restrepo (2018d) and empirically estimate its parameters (see Arntz et al., 2019b for details). In particular, their model covers the following mechanisms:

1. *Substitution and complementarity*: Technologies both substitute for some workers while complementing others. They substitute for workers, as the automation technologies replace workers in tasks they previously performed (*displacement effect*). They simultaneously complement other workers, since they require more input of other types of workers who do tasks complementary to the machines. These could be, for example, new tasks, in which case this effect is similar to the *reinstatement effect*, above. Given a certain output level, firms' investments thus reduce demand for the former workers, while raising demand for the latter. Arntz et al. (2018) estimate the net effect of substitution and complementarity.
2. *Product demand*: (a) Technology investments affect firms' competitiveness, which reduces prices and increases output, thereby raising labor demand and employment. (b) In addition, technological capital has to be produced, which implies that any change in firms' investment decisions affects employment via capital production. (c) Finally, the expanding technological frontier implies that the economy can produce more as a whole and becomes richer, which raises consumption, production, and employment. These effects are similar to the *productivity effect*, discussed above.
3. *Labor supply*: Changes in the demand for labor affect unemployment which induces wage responses. Both, in turn, trigger worker mobility from declining labor market segments to growing labor market segments. The resulting effects are ambiguous: On the one hand, occupations and industries with declining demand will experience falling wages, thereby reducing the cost incentives for automation and thus limiting the employment consequences. At the same time, workers will try to leave these segments, which reduces the increase in the segments' unemployment rate, hence limiting the described wage response. Labor supply responses may thus either limit or amplify the employment consequences of technology-induced labor demand shocks.

Arntz et al. (2019b) estimate the parameters of this model using the previously mentioned novel LM4.0 firm survey, employment data from social security records, international trade data, and official statistics. That is, they exploit the fact that the LM4.0 survey covers firms that already invested in cutting-edge technologies between 2011 and 2016 in order to study how the economy and the labor market respond to the adoption of new technologies. Arntz et al. (2018) then feed this model with firms' investment plans for the next five years to study how these investment plans are likely to affect the German labor market until 2021.

4.2 Baseline Scenario

This section discusses the likely consequences of firms' investment plans for employment and wages in Germany in the next five years. The model allows to empirically disentangle three mechanisms by which firms' technology investments affect employment, as outlined above. Figure 7 presents the main results from the baseline simulation. Overall, firms' future plans for




Source: Arntz et al. (2018)

Figure 7: Employment Effects

technology investments are expected to increase employment by 1.8% over five years. Thus, contrary to public fears, new technologies might actually raise rather than reduce employment.

Figure 7 plots the effects of firms' investments into all technology types on employment via the different mechanisms. That is the calculations take into account investments into all technology types, both recent and older technologies, and at first do not disentangle the effects by the different generations of technologies. Intuitively, one would expect the substitution effect to dominate the complementarity effect when automation technologies are introduced. Quite unexpectedly, the opposite holds true. The introduction of the new technologies seems to require more additional labor input in complementary occupations than it can substitute in replaceable occupations. The new technologies require more rather than less workers. The overall net positive employment effect therefore does not stem from positive product demand effects but rather from the fact that new technologies seem to be complementary rather than substitutable to work. Quite the opposite, the product demand effect is actually negative indicating that firms cannot gain from lower costs and demand expansions. This suggests that firms currently invest in the adoption of the new technologies which requires more workers than can be replaced by machines and where high investment costs at least temporarily imply that firms cannot gain from lower costs and demand expansions. However, a positive effect on employment remains, which puts upward pressure on wages. The resulting wage increase limits the employment expansion – this is reflected in a negative contribution of the labor supply effect to overall employment.

Next, Figure 8 differentiates by type of technology in order to disentangle how these types differently affect workers. The types of technologies are explained in Section 3.1. The pattern for older, more mature technologies strikingly differ from cutting edge 4.0-technologies. While for technologies up to technology-level 3.0 substitution effects dominate complementary effects, the opposite is true for 4.0-technologies. Investments in older technologies initially replace more workers than they require, but induce positive product demand effects as firms get more productive. This limits negative net employment effects. In contrast, the complementarity effects of most recent technologies clearly dominate their substitution potential and induce negative rather than positive product demand effects. This suggests that firms currently hire



Source: Arntz et al. (2018)

Figure 8: Employment Effects by Technology Type

workers to adopt these technologies without being able to substitute for lots of workers, so far. Therefore, firms which invest into these technologies cannot yet achieve large efficiency gains, such that they lose competitive advantages in the short run, resulting in lower output and less employment. Yet, the net effect on employment of the newest technologies remains positive.

These patterns observed for 4.0 technologies match the description of Brynjolfsson et al. (2019) of such technologies being a GPT that first requires complementary investments before unfolding its productivity potentials in the longer run (see Section 3.1). This would also suggest that the employment effects of 4.0 technologies are likely to change once they start to mature. Hence, to the contrary of widespread beliefs, it is currently not the investments in the newest technologies that substitute for workers, but rather investments in older technologies that are less likely to go along with the creation of new tasks. Moreover, these findings support the view that the productivity puzzle, i.e. the stagnant labor productivity despite high investments in new technologies, may actually reflect that we are experiencing an investment period whose returns also in terms of rising labor productivity likely still need time to unfold. While the results from this section support this explanation for the productivity puzzle, they do not rule out competing explanations. Others, for example, argue that these technologies actually do not have large productivity effects (e.g. Gordon, 2014, 2015), that innovation and technological progress are slowing down (e.g. Cowen, 2011; Bloom et al., 2017), that the focus of technological change has shifted towards automation that has little productivity effects (e.g. Acemoglu and Restrepo, 2018a), or that measurement errors prevent observing the productivity effects (e.g. Mokyr, 2014).

Despite the fact that the model predicts no net employment losses of technology investments, it does predict strong structural changes on the labor market such as a reallocation of employment across occupations and sectors. To demonstrate this, Figure 9 reports the joint employment effects of investments in all types of technology by type of occupation. As expected, the effects are particularly positive for workers in occupations with a focus on analytical and interactive tasks. These are tasks that are unlikely to be substituted by new technologies, but which instead are complementary to these. The picture is very different for workers in occupations that demand high shares of cognitive routine-tasks. These tasks are most exposed to



Source: Arntz et al. (2018)

Figure 9: Employment Effects by Occupation

automation via new technologies and also unlikely to benefit from newly created tasks. As a consequence, workers focusing on such tasks suffer from technological change. Interestingly, the effects on mainly manual occupations is rather small, emphasizing that current technological advances mainly affect cognitive tasks.

Figure 10 plots the employment effects by the initial average daily wage for each segment of the labor market. There are in total 60 labor market segments, resulting from 5 occupational groups and 12 industry aggregates. Cells with high initial average daily wages benefit most from technology investments and expand, whereas low- and particularly medium-paid occupations and industries face stagnating or even declining employment. Firms' planned technology investments thus are expected to induce rising inequality and a (weak) employment polarization.

4.3 Moderating Factors

In order to study how wage setting frictions and workers' mobility moderate the effect of technological change on the labor market, Arntz et al. (2018) develop two scenarios. In the *mobility scenario*, they simulate employment and wage responses to technological change assuming the mobility elasticity to be twice as high as in the baseline model. In the *rigid wages scenario*, they conduct a similar analysis assuming the wage elasticity to be half as large as in the baseline model.

The main employment effects of these additional scenarios can be depicted from Figure 11. Obviously, when solely changing elasticities which enter the labor supply part of the framework while leaving labor and product demand unchanged, only the labor supply effect changes. In the mobility scenario, hardly any change in the labor supply effect is visible. This is due to the fact that the size of the overall workforce is fixed in the model. Workers now move faster between labor market segments thus changing the allocation of workers across occupations and sectors, but this hardly affects the overall net labor supply effect. In the rigid wages scenario, the negative labor supply effect is smaller instead. As wages are rigid, they do not rise as fast in the expanding labor market segments, such that the expansion of employment in these segments is less limited by wage increases. Overall, rigid wages amplify any effects that technologies have on






Figure 10: Employment Effects by Occupation



Source: Arntz et al. (2018)

Figure 11: Employment Effects for Different Scenarios



Source: Arntz et al. (2018)

Figure 12: Wage Effects for Different Scenarios

employment. In the present case, rigid wages thus results in larger positive effects of technology on net employment, since a smaller part of the net positive labor demand shock accrues to wages, and a larger part accrues to employment.

The flip-side of rigid wages thus is slower wage growth compared to the baseline scenario, as outlined in Figure 12. This holds across all occupations. Nevertheless, the fast expanding analytical and interactive occupations still face higher wage growth. In the mobility scenario, differences in wage growth between occupations are less pronounced. This is due to the fact that workers are more mobile between occupations and switch faster to the expanding occupations, thereby limiting wage growth in these occupations. Simultaneously, labor supply becomes scarcer in the declining occupations, raising wages in these occupations compared to the baseline scenario. Hence, the higher worker mobility rates not only help the mobile workers to achieve higher wages in other occupations, but they also benefit those workers who remain in the declining occupations by reducing competition between them. Thus, even if higher mobility of workers has hardly any effect on overall employment, it helps a larger share of workers to reap the benefits of technological change and, hence, also reduces wage and employment inequality and polarization.

4.4 Implications

The results from these simulations provide five key results for the likely effects of automation and digitalization on the German labor market in the next five years:

(1) Firms' plans to invest in automation and digitalization technologies likely have small positive effects on employment in Germany. In the baseline scenario, these investments raise overall employment by 1.8% in 5 years. Mass technological unemployment thus remains unlikely, a finding that is in line with the impact of technological change in earlier decades (see Section 3.3). However, these net positive employment effects do not imply an absence of job losses. In the scenarios, automation does destroy jobs in specific occupations and industries, as automation technologies replace workers. However, these negative effects are more than compensated by job creation effects of automation, resulting in overall net positive employment effects.

(2) The small net positive employment effects are accompanied by large structural shifts between occupations and industries in the 5-year scenario. Hence, the key challenge of automation and digitalization does not relate to the number of jobs, but to the job structure. In line with much of the literature, cognitive routine jobs continue to decline, as they are replaced by machines. Jobs with high shares of abstract or interactive tasks, on the other hand, are on the rise, as these are typically complementary to the new technologies. Workers who are able to switch to these jobs therefore are likely to gain from the adoption of new technologies, whereas workers who lack such skills will likely suffer from automation and digitalization.

(3) The expanding abstract and interactive task intensive occupations typically are high-wage occupations, whereas the stagnating occupations are located in the middle and at the lower end of the wage distribution. This implies that mostly high-skilled and well-paid workers profit the most from digitalization, whereas middle- and lower-skilled workers fall further behind. Digitalization and automation, therefore, likely raise inequality in Germany. The results point to a continued employment polarization, as the jobs in the middle seem to grow even slower than those at the lower tail, although rising inequality dominates. There exists a similar pattern of weak wage polarization accompanied by markedly rising wage inequality in response to technological change. Hence, in line with results from studies on automation potentials (see section 2), technological change is likely to raise inequality on the German labor market in the next five years by favoring high-skilled workers.

(4) Our simulations further highlight that rising worker mobility, e.g. via training and further education, can help to mitigate rising inequality and to ensure that a larger fraction of workers profits from technological change. While the simulations suggest that rising mobility has hardly any effect on overall employment or unemployment, it is still beneficial to many workers by either enabling them to take up better-paid jobs in expanding occupations, or by reducing the pressure in declining occupations.

(5) Finally, the simulations strongly suggest that automation and digitalization at least currently is a costly investment for many firms. In particular, firms that invest in cutting-edge technologies face a rising demand for certain types of workers, while so far not being able to replace workers on a large scale. In addition, there appears to be no or only little expansion in the output of those firms. All this suggests that firms currently incur high investment costs to adopt new technologies, while not being able to reap related benefits in terms of higher productivity and lower costs, yet. Once this investment phase is completed, the new technologies may unfold their productivity advantages. This phase may then be accompanied by larger technology-induced job-separations, as one actually observes for investments into older, more mature technologies. However, the related productivity gains simultaneously induce an increasing product demand and likely result in the generation of new jobs.

These key results entail three main implications for policy makers.

Promoting new technologies: Overall, technological change contributes to employment growth. In the case of investments in “3.0 technologies”, this is mainly the case due to the strongly positive productivity effects, whereas cutting-edge “4.0 technologies” currently seem to complement rather than replace workers due to their investment character. Corresponding investments thus require significantly more skilled workers for the implementation of these technologies. An accelerated diffusion of both 3.0 and 4.0 technologies into companies therefore may be a desirable goal as, on net, both technologies raise employment. Policy measures to support the adoption of new technologies (e.g. broadband expansion, data protection laws) could thus help to increase related gains. In particular, the results based on the LM4.0 Survey indicate that the technological latecomers seem to lack information to better assess the opportunities and risks related to these technologies. Targeted information campaigns, e.g. at the level of industry associations and regionally organized networks, may help to reduce information deficits.

Addressing shortages of skilled workers: In the medium term, new digital technologies are strongly complementary to analytical and interactive activities. The growth potential resulting from the new technologies thus depends strongly on the availability of suitably skilled workers. Here, appropriate educational policies help to ensure that the skills in demand are trained both in schools and in the area of vocational and university education. In addition, the number of skilled workers may also be increased by further training measures. However, which measures are most likely to ease occupational transitions can hardly be derived at the aggregate level, but requires further analysis at the individual level.

Increasing mobility: The results on labor supply show that mobility between labor segments is currently relatively high and that increased labor mobility has little impact on the net employment effects. However, increased mobility between shrinking and growing labor market segments contributes to counteracting employment and wage inequality. Accelerated mobility from shrinking segments to growing segments leads to an alignment of employment opportunities and wage developments in the segments. Training and qualification measures thus seem natural recommendations to raise mobility. Nevertheless, in order to make more targeted recommendations on how mobility between different occupations and sectors can be increased, further analysis at the individual level are needed, to analyze, for example, the influence of further training and qualification measures.

5 Conclusions

The past decades have been characterized by a tremendous rise in computing power, reducing the costs of automating so-called routine tasks which follow clear, explicit rules and can thus be put into computer code. This has led to a polarization of labor markets in advanced economies with declining shares of middle-paid, routine-intensive occupations and rising shares of both, high- and low-paid jobs.

While this computerization has not led to employment declines, the question whether this holds true for the effects of further technological advances in the near future remains open. Whereas previous automation methods were limited to problems that are sufficiently well understood to be put into algorithms of well-defined steps, now even less structured problems appear automatable using big data and machine learning. Continued increases in computing power, the growing availability of big data, and significant advances in Machine Learning methods are shifting the boundaries of what can be automated by machines. Against this background, some studies predict that about half of the U.S. workforce is “at risk of automation”, which has spurred public fears of technology-induced mass unemployment.

This chapter contrasts such fears with the scientific debate. The first main contribution is to show that many estimates of automation potentials are severely upward biased, as they often are conducted at the occupational level, ignoring the huge heterogeneity of what people actually do at work. As many workers in seemingly automatable occupations already adjust their task schedules to non-automatable tasks, they often face much lower exposure to automation. This chapter finds that the share of workers in automatable jobs is more in the order of 9% in the U.S., and similarly in other countries.

These numbers, however, only refer to technological potentials and must not be equated with actual job losses or employment effects as is often done in the public debate. The second main contribution of this chapter is to explain why this is the case. In particular, there are three main reasons for this: (1) The diffusion of new technologies into the economy is a rather slow process, leaving workers time to adjust. Diffusion is slow due to high costs, uncertainty, the need to undergo organizational change for implementing the technologies, and the need for

acquiring workers with suitable skills. (2) Workers are flexible and adjust. In fact, much of the adjustment to automation is not made by making seemingly replaceable occupations redundant, but by workers doing other tasks in the same occupations. Being in an occupation that is “at risk” thus does not necessarily imply that the worker is about to lose his or her job, but that the worker has to adapt by switching to the right tasks and learning the right skills. (3) Finally, while automation indeed does displace jobs, it simultaneously creates new jobs. The overall effect on the number of jobs (employment) has been actually positive, not negative. It is thus ambiguous, whether the wave of new automation technologies will reduce or actually raise labor demand.

Whether the next wave of digitalization and automation thus leads to less or more jobs, is an open question. The third contribution of this chapter is to present scenarios for the potential impact of digitalization and automation via cutting-edge technologies on the German labor market, exploiting a recent survey on the adoption of new digital technologies and a new framework to estimate and simulate the effects. The results suggest that the net effect remains small, and is actually positive in the next five years. However, there appear large structural shifts between occupations and industries, which are accompanied by rising inequality and, weakly, by employment polarization. The main challenge for the future thus is not mass unemployment, but structural change. In addition, the simulations suggest that we currently experience an investment phase, where firms first have to incur high investment costs and need to acquire the right skilled workers, before being able to reap large productivity gains. Hence, the effects of these cutting-edge technologies may change in the medium to long-run, when the technologies mature. Nevertheless, this does not imply that they reduce employment in the longer run, as, once they mature, they simultaneously create productivity effects that also raise demand for labor. It remains to be seen whether the job-creating effects continue to dominate the job-destruction effects in the longer run.

These results entail three main policy implications. Firstly, promoting the adoption of new technologies seems to be a reasonable policy goal, as these technologies apparently raise employment and production. The focus should be on medium and small firms who currently seem to fall behind. Secondly, the introduction of these technologies requires workers with the right skills. The lack of such workers seems to partly hinder the introduction of new technologies. The second recommendation thus is to address skill shortages by education, qualification, and further training. Finally, the coming wave of technological change seems to be associated with a further rise in inequality, as high-skilled, high-wage occupations are on the rise, whereas low- and medium paid jobs further fall behind. In order to prevent further rising inequality, targeted training and qualification measures may help workers to switch to the expanding occupations, thus helping them to participate in the technology-induced benefits, while lowering the losses of those who cannot change their skills and jobs and thus remain in shrinking occupations and sectors.

References

- Acemoglu, D. (1998). Why do new technologies complement skills? directed technical change and wage inequality. *Quarterly Journal of Economics*, 113(4):1055–1089.
- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Card, D. and Ashenfelter, O., editors, *Handbook of Labor Economics*, chapter 12, pages 1043–1171. Elsevier.
- Acemoglu, D. and Restrepo, P. (2017). Robots and jobs: Evidence from us labor markets. NBER Working Paper No. 23285.
- Acemoglu, D. and Restrepo, P. (2018a). Automation and new tasks: The implications of the task content of production for labor demand. *Journal of Economic Perspectives*, forthcoming.
- Acemoglu, D. and Restrepo, P. (2018b). Low-skill and high-skill automation. *Journal of Human Capital*, 12(2):204–232.
- Acemoglu, D. and Restrepo, P. (2018c). Modeling automation. *American Economic Review: Papers & Proceedings*, 108:48–53.
- Acemoglu, D. and Restrepo, P. (2018d). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.
- Arnold, D., Arntz, M., Gregory, T., Steffes, S., and Zierahn, U. (2016). Herausforderungen der digitalisierung für die zukunft der arbeitswelt. ZEW policy brief Nr. 16-08.
- Arntz, M., Genz, S., Gregory, T., Janser, M., Lehmer, F., Matthes, B., and Zierahn, U. (2019a). Technology and jobs in the fourth industrial revolution - firm-level evidence.
- Arntz, M., Gregory, T., and Zierahn, U. (2016). The risk of automation for jobs in oecd countries: A comparative analysis. OECD Social, Employment and Migration Working Papers No. 189.
- Arntz, M., Gregory, T., and Zierahn, U. (2017). Revisiting the risk of automation. *Economics Letters*, 159:157–160.
- Arntz, M., Gregory, T., and Zierahn, U. (2018). Digitalisierung und die zukunft der arbeit: Makroökonomische auswirkungen auf beschäftigung, arbeitslosigkeit und löhne von morgen. Technical report, Bundesministerium für Forschung und Entwicklung (BMBF).
- Arntz, M., Gregory, T., and Zierahn, U. (2019b). Automation, employment, and wages: Is this time different?
- Autor, D. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.
- Autor, D. and Dorn, D. (2013). The growth of low skill service jobs and the polarization of the u.s. labor market. *American Economic Review*, 29(3):1553–1597.
- Autor, D. and Handel, M. (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of Labor Economics*, 31(2):S59–96.
- Autor, D., Levy, F., and Murnane, R. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.

- Bessen, J. (2015). Toil and technology. *Finance & Development*, 2015(March):16–19.
- Bessen, J., Goos, M., Salomons, A., and van den Berge, W. (2019). Automatic reaction - what happens to workers at firms that automat? Law & Economics Series Paper No. 19-2, Boston University.
- Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2017). Are ideas getting harder to find? NBER Working Paper No. 23782.
- Bonnefon, J., Shariff, A., and Rahwan, I. (2016). Autonomous vehicles need experimental ethics: Are we ready for utilitarian cars? *Science*, 352:1573–1576.
- Bowles, J. (2014). The computerisation of european jobs. Bruegel Blog, 17th July 2014.
- Bresnahan, T., Brynjolfsson, E., and Hitt, L. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, 117(1):339–376.
- Brynjolfsson, E. and McAfee, A. (2016). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. Norton & Company.
- Brynjolfsson, E. and Mitchell, T. (2017). What can machine learning do? workforce implications. *Science*, 358(6379):1530–1534.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In Agrawal, A. K., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, chapter 1.
- Cortes, M. (2016). Where have the middle-wage workers gone? a study of polarization using panel data. *Journal of Labor Economics*, 34(1):63–105.
- Cortes, M. and Salvatori, A. (2018). Delving into the demand side: Changes in workplace specialization and job polarization. *Labour Economics*, 57:164–176.
- Cowen, T. (2011). *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*. Dutton.
- Dauth, W., Südekum, J., and Wössner, N. (2017). German robots - the impact of industrial robots on workers. IAB Discussion Paper 30/2017.
- Dengler, K. and Matthes, B. (2018). The impacts of digital transformation on the labour market: Substitution potentials of occupations in germany. *Technological Forecasting & Social Change*, 137:304–316.
- Feldmann, H. (2013). Technological unemployment in industrial countries. *Journal of Evolutionary Economics*, 23:1099–1126.
- Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerization? *Technological Forecasting & Social Change*, 114:254–280.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- Gordon, R. J. (2014). The demise of us economic growth: Restatement, rebuttal, and reflections. NBER Working Paper No. 19895.

- Gordon, R. J. (2015). Secular stagnation: A supply-side view. *American Economic Review*, 105(5):54–59.
- Graetz, G. and Michaels, G. (2018). Robots and work. *Review of Economics and Statistics*, 100(5):753–768.
- Gregory, T., Salomons, A., and Zierahn, U. (2018). Racing with or against the machine? evidence from europe. CESifo Working Paper No. 7247.
- Janssen, S. and Mohrenweiser, J. (2018). The shelf life of incumbent workers during accelerating technological change: Evidence from a training regulation reform. IZA Discussion Paper No. 11312.
- Jovanovic, B. and Rousseau, P. L. (2005). General purpose technologies. In Aghion, P. and Durlauf, S. N., editors, *Handbook of Economic Growth*, chapter 18.
- Mokyr, J. (2014). Secular stagnation? not in your live. In Teulings, C. and Baldwin, R., editors, *Secular Stagnation: Facts, Causes and Cures*, pages 83–89. CEPR Press.
- Mokyr, J., Vickers, C., and Ziebarth, N. L. (2015). The history of technological anxiety and the future of economic growth: Is this time different? *Journal of Economic Perspectives*, 29(3):31–50.
- Nedelkoska, L. and Quintini, G. (2018). Automation, skill use and training. OECD Social, Employment and Migration Working Papers No. 202.
- Nordhaus, W. D. (2007). Two centuries of productivity growth in computing. *Journal of Economic History*, 67(1):128–159.
- Pajarinen, M. and Rouvinen, P. (2014). Computerization threatens one third of finnish employment. ETLA Brief 22, January 2014.
- Pianta, M. (2009). Innovation and employment. In Fagerberg, J., Mowery, D., and Nelson, R., editors, *The Oxford Handbook of Innovation*, chapter 21.
- Pouliakas, K. (2018). Determinants of automation risk in the eu labour market: A skills-needs approach. IZA Discussion Paper No. 11829.
- Pratt, G. A. (2015). Is a cambrian explosion coming for robotics? *Journal of Economic Perspectives*, 29(3):51–60.
- PWC (2018). Will robots really steal our jobs? an international analysis of the potential long term impact of automation. PriceWaterhouseCoopers.
- Rifkin, J. (1995). *The End of Work: The Decline of the Global Labor Force and the Dawn of the Post-Market Era*. Putnam Publishing Group.
- Solow, R. (1987). We’d better watch out. New York Times Book Review, July 12 1987, page 36.
- Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics*, 24(2):235–270.
- Thierer, A. and Hagemann, R. (2015). Removing roadblocks to intelligent vehicles and driverless cars. *Wake Forest Journal of Law & Policy*, 5(2):339–391.
- Vivarelli, M. (2007). Innovation and employment: A survey. In Hanusch, H. and Pyka, A., editors, *Elgar Companion to Neo-Schumpeterian Economics*, pages 719–32. Elgar.