



Editorial

The governance of artificial intelligence: Harnessing opportunities and mitigating challenges[☆]

1. Introduction

1.1. Definition of AI

The OECD defines an AI (Artificial Intelligence) system as “a machine-based system that can influence the environment by producing an output (predictions, recommendations, or decisions) for a given set of objectives. It uses machine and/or human-based data and inputs to (i) perceive real and/or virtual environments; (ii) insert these perceptions into models through analysis in an automated manner (e.g., with machine learning), or manually; and (iii) use model inference to formulate options for outcomes. AI systems are designed to operate with varying levels of autonomy” (OECD, 2019a).

This definition relates AI to the type of technology that has created the recent excitement around technological progress: machine learning. Machine learning is a branch of computational statistics that focuses on designing algorithms to make predictions from new data without explicitly programming the solution. Since 2012, the use of machine learning as a prediction technology has grown substantially. Machine learning is now commonplace: Pandora learns how to make better music recommendations based on its users’ preferences; Google learns how to automatically translate content into different languages based on translated documents found online; and Facebook learns how to identify people in photos based on its database of known users.

One set of machine-learning algorithms called “deep learning” has proven to be particularly useful and commercially viable for a variety of prediction tasks. Deep-learning algorithms are neural networks that solve problems from large, complex datasets with very little guidance from programmers. A neural network is a program that uses a combination of weights and thresholds to translate a set of data inputs into predictions for outputs, measures the “closeness” of these predictions to reality, and then adjusts the weights it uses to narrow the distance between predictions and reality. In this way, a neural network can learn as it is fed more data. It is called “deep” learning because the program automatically generates multiple nets as layers of abstractions of the data to identify patterns.¹

While recent interest in AI is driven by machine learning, computer scientists and philosophers have emphasized the feasibility of a true artificial general intelligence that equals or exceeds human intelligence,

also known as AGI (Artificial General Intelligence). The promise of AGI has been around at least since the 1950s. The first AI conference was held at Dartmouth College in 1956. In 1965, Nobel laureate Herbert Simon famously quipped that “machines will be capable, within 20 years of doing any work a man can do.” However, progress in AGI has come slowly, including the distinct period of an “AGI winter” from 1974 to 1987, when progress slowed, and funding decreased. In recent years, interest in AGI has seen a resurgence in the wake of developments in machine learning. More recently, a Large Language Model (LLM) has been launched by Microsoft’s OpenAI. ChatGPT (Generative Pre-Trained Transformer) is autoregressive and uses deep learning to reproduce and refine human text on demand. This is the closest technological progress has come to AGI, though it has raised new ethical challenges, particularly concerning the boundaries (and relationship) between humans and machines.

1.2. Opportunities and challenges coming from AI

As AI continues to evolve and find its way into a wide variety of applications, such as the proliferation of new and free goods including search engines or social media, AI will generate welfare. Currently, however, AI’s economic value is not well-captured in our current national accounts. To this end, Brynjolfsson et al. (2019) propose a new metric called GDP—B, which quantifies the benefits of digital goods and services rather than their costs. Through a series of choice experiments, they estimate consumers’ willingness-to-pay for free digital goods and services. For example, the welfare gains from Facebook would have added between 0.05 and 0.11 percentage points per year to growth in US GDP. These are significant changes, especially considering that Facebook is just one product in the digital economy.

However, for society to benefit from the impact of AI, governance of AI is essential. Acemoglu (2021) argues that only by carefully managing the development and impact of AI will society benefit from it. He identifies several challenges. One such challenge is the collection and control of data. Examples include privacy violation, unfair competition through unequal access to data, and behavioral manipulation by machine-learning algorithms that enable companies to identify and exploit biases and vulnerabilities among consumers. Another challenge is related to the impact of AI on communication in society and

[☆] This paper is the Introduction to the Special Issue on The Governance of AI.

¹ Dean (2019) summarizes the evolution of deep learning. Key ideas and algorithms underlying deep learning have been around since the 1960s. At that time, though, computers were not powerful enough to allow this approach to work on anything but small problems. It was not until around 2012, after decades of computational performance improvements driven by Moore’s Law, that computers finally started to become powerful enough to train large neural networks on realistic, real-world problems.

democracy. Of particular concern here are echo chambers in social media that propagate false information and polarize society, or the ability of governments to closely monitor dissent through AI. The final challenge is the impact of AI on work, including AI's potential to automate instead of augment workers or to monopolize worker information and excessively monitor workers.

Arguably, another significant challenge for the future of innovation and the generation of original knowledge is the regulation of the boundary between human and artificial intellectual property rights. Recently, the Authors Guild, a prominent organization representing writers in the US, has initiated a class-action lawsuit against OpenAI, claiming, among other things, direct copyright infringement for using code, intellectual and artistic work to train generative AI, in addition to unfair competition and unjust enrichment for not paying copyright fees. While the US Court of Justice seems so far inclined to protect human authorship as a 'bedrock requirement of copyright', therefore ruling out that AI-generated art can be copyrighted, this might become a slippery slope.

AI's potential cannot be realized without a proper understanding and management of these challenges. Mokyř (2005) further highlights the importance of this point in a historical perspective. He argues that sustained economic growth after the onset of the Industrial Revolution around 1750 was not solely due to the specific inventions made, but also facilitated by our understanding and management of these inventions. While there was also economic growth before 1750 from inventions such as gunpowder, spectacles, or the mechanical clock, this growth was not sustained because of a lack of understanding and management of these technologies. This changed after 1750 when markets and institutions became more supportive of sustaining inventions from the Industrial Revolution such as steam power, electricity, combustion engines, railroads and air travel, or telephones. Similarly, over the past four decades scientific research has progressed to a better understanding of digital computing and its impact on society. But looking ahead, how much do we understand the societal consequences of AI? And what are the main directions lying ahead for the governance of AI to make sure that these consequences are beneficial, sustainable and just?

This special issue aims to contribute to answering these questions. The remainder of this article summarizes the articles within this special issue, while embedding them in the recent literature. Section I discusses the recent progress in our knowledge of the development and adoption of AI. Section II turns to the impact of AI on work. Finally, Section III focusses on the governance of data.

2. Development and adoption of AI

2.1. Development of AI

The recent development of AI can have a large impact on the economy by serving as a new General-Purpose Technology (GPT). Cockburn et al. (2019) argue that deep learning may have an even larger impact by also serving as an "Invention in the Method of Invention" (IMI). What sets IMIs apart from GPTs is that IMIs can also reshape the nature of the innovation process and the organization of R&D itself. On the one hand, deep learning may be able to substantially "automate discovery" across many domains where classification and prediction tasks play an important role. On the other, deep learning may also "expand the playbook" in the sense of presenting the set of problems that can be feasibly addressed.

Previous IMIs help illustrate their importance. For example, the invention of optical lenses had important and direct economic impact on applications such as spectacles. But optical lenses in the form of microscopes, invented in the 17th century, also had enormous and long-lasting indirect effects on the progress of science: by making very small objects visible for the first time, microscopes created the field of micro-biology. Today, deep learning enables us to better understand genomes, thereby progressing the fields of molecular biology and

genetics.

In this special issue, Bianchini et al. (2022) look at the adoption of Neural Network and AI in the domain of scientific research and extend the IMIs by proposing the concept of *emerging general method of invention*. The role of AI in scientific research is empirically shown to align to the characteristics of an *emerging technology*, that is:

"[a] radically novel and relatively fast-growing technology characterized by a certain degree of coherence persisting over time and with the potential to exert a considerable impact on the socio-economic domain(s) which is observed in terms of the composition of actors, institutions and patterns of interactions among those, along with the associated knowledge production processes. Its most prominent impact, however, lies in the future and so in the emergence phase is still somewhat uncertain and ambiguous" _ (p.1828). (Rotolo et al., 2015)

However, Bianchini et al. (2022) also find that AI has not yet been able to expand the re-combinatorial element of scientific discovery that makes it possible for scientific research to create novelty by cross-fertilizing (very) diverse scientific domains, although it is already displaying its potential to expand scientific discovery within each scientific domain. In other words, AI in science is augmenting the intensive margin but not yet the extensive margin of knowledge production.

In this context, policy and institutional responses will also be required if deep learning is to represent a meaningful IMI or an emerging general method of invention. An important policy challenge is that control, both in the form of physical exclusivity as well as in the form of formal intellectual property rights over tools and data, can shape both the level and direction of innovative activity. If there are increasing returns to scale or scope in data acquisition for researchers, it is possible that some will gain a long-lasting advantage over others merely through the control of data. Moreover, strong incentives for researchers to maintain data has the potential downside that data will not be shared, thus reducing the ability of all researchers to access large sets of data that would arise from public aggregation. We will explore these issues further in Section III.

The importance of public policy in AI innovation is illustrated in this special issue by Mateos Garcia et al. (2024), who focus on the overall direction of AI research. Their starting point is that a myopic focus on short-term benefits could limit AI to technologies that turn out to be sub-optimal in the longer run. For this reason, it may be useful to preserve diversity in the AI trajectories. Using arXiv, a widely used pre-prints site, they identify 110,000 AI papers to estimate the thematic diversity of AI research. Their results show that diversity in AI research has stagnated in recent years, and that AI research involving private sector organisations tends to be less diverse than research in academia. This appears to be driven by a small number of prolific and narrowly focused technology companies, while diversity in academia is bolstered by smaller institutions and research groups. They also find that private sector AI researchers tend to specialize in data and computationally intensive deep-learning methods at the expense of research involving other AI methods and of research that considers the societal and ethical implications of AI or applies it in domains like health. In sum, their results suggest that there may be a rationale for policy action to prevent the premature narrowing of AI research that could reduce its societal benefits, despite the incentive, information and scale hurdles standing in the way of such interventions.

2.2. Adoption of AI

Acemoglu et al. (2022a) provide a comprehensive description of the adoption of automation technologies by US firms across all economic sectors by leveraging a new module introduced in the Census Bureau's 2019 Annual Business Survey (ABS). The module collects data from over 300,000 firms on the use of five advanced technologies: AI, robotics,

dedicated equipment, specialized software, and cloud computing (which is often used in combination with one of the other four advanced technologies). Overall, they find that the adoption of AI and robotics has remained relatively low. In particular, only 3.2 % of US firms currently use AI as part of their processes and methods and 2 % use robotics. Instead, 19.6 % of firms use dedicated equipment, 40.2 % use specialized software, and 34 % use cloud computing. Still, half of US firms had used none of these advanced technologies by 2018.

They also find considerable differences in the adoption of advanced technologies between sectors. For example, robotics remains highly concentrated in manufacturing, with 9 % of manufacturing firms using robots. Because these robot adopters are mainly large firms, 45 % of all manufacturing workers are exposed to this technology. Specialized software and cloud computing are used most, with penetration rates varying between 20 % and 60 % across all sectors of the economy. The use of dedicated equipment is also pervasive across sectors, but the fraction of firms in each sector using this technology is somewhat lower on average with a minimum of 10 % (in Finance, Insurance and Real Estate) and a maximum of 40 % (in Manufacturing). AI is also being adopted in all sectors of the economy, but adoption rates are still low on average, with the highest rates observed for Information (12 %), Professional Services (10 %), Finance, Insurance and Real Estate (7 %), and Health Care (5 %).

Important differences in AI adoption also exist within sectors. First, larger firms are more likely to adopt AI. For example, 2.5 % of firms with employment below the median have adopted AI, in comparison to 5 % of the top 1 % largest firms (in terms of employment). Second, irrespective of firm's size, younger firms are much more likely to adopt AI. For example, of all large firms in the 95th to 99th percentile of the firm size distribution, 7 % of firms in the youngest age-quartile have adopted AI, whereas only 3.5 % of firms in the oldest age-quartile have. The fact that AI adoption concentrates in larger and younger firms reflects that there are substantial costs and organizational barriers involved in adopting AI. The ABS module contains some questions to further examine these barriers to AI adoption, suggesting that the inapplicability of AI to the firm's business and AI being too costly are the main reasons for not adopting AI.

The ABS module also asks why firms adopt AI. Of all AI adopters, 80 % (employment-weighted) report to have done so to improve the quality of their product or service, 65 % to upgrade existing processes, and 55 % to automate existing processes. The use of AI to automate existing processes could have important adverse consequences for workers (we will return to this question in the next section). Importantly, this use of AI is different from the reasons firms have in adopting other advanced technologies: only 37 % of firms adopting dedicated equipment, 32 % of firms using specialized software, and 23 % of firms using cloud computing do so for the purpose of automating existing processes. The only exception is robotics, with 65 % of robot-adopters doing so to automate the workplace. In sum, just like robotics, current AI competes more intensively with workers than other advanced technologies.

Similar insights exist for European countries, although survey designs and taxonomies of AI applications differ from the ABS module for the US. In 2020, Eurostat surveyed enterprises on their use of AI applications. The survey excludes financial sector companies and micro-enterprises with fewer than 10 employees. Eurostat asked about the following four types of AI: those that analyze big data using machine learning; those that analyze big data using natural language processing; those that use chatbots; and those that use service robots. In 2020, 7 % of enterprises in the EU reported to be using one of those four types of AI. While 2 % of the enterprises used machine learning to analyze big data, 1 % did so with the help of natural language processing. A chat service, in which a chatbot or virtual agent generated natural language replies to customers, was used in 2 % of the enterprises. The same proportion of enterprises, 2 %, used service robots.

Among the EU Member States, Ireland recorded the highest share of enterprises (23 %) that used any of the four considered AI applications in

2020. Other countries with widespread uptake of AI technologies were Malta (19 %), Finland (12 %) and Denmark (11 %). In contrast, less than 10 % of enterprises used any of the four AI applications in 2020 in all other Member States. The lowest shares were recorded in Latvia (2 %), Slovenia, Hungary, Cyprus (3 % each) and Poland (4 %).

Hoffmann and Nurski (2021) discuss the Eurostat as well as other surveys (including a less representative survey by the European Commission) in more detail. As in the US, they find that robots are concentrated in manufacturing, while the adoption of other types of advanced technologies is higher in services such as finance, education, health, and social work. Within each of these sectors, larger firms are more likely to adopt AI, suggesting that there are substantial costs and organizational barriers involved in adopting AI also in Europe. Skills and financial constraints are the leading reported barriers, with about 80 % of enterprises citing a lack of skills in their internal workforce and in the external labour market, as well as the high cost of buying the technology and adapting their operational processes to AI.

More recently, Calvino and Fontanelli (2023) look at harmonized microdata from the AI diffuse project at the OECD² and find that AI adoption is more widely diffused in large, young and professional services firms, with less financial barriers and more capacity for complementary assets such as digital infrastructure and digital skills of the workforce.

What are the enabling factors and complementary assets that firms rely upon for AI adoption to positively impact firm growth? On the basis of unique survey data of AI in (US) startups, in this special issue Bessen et al. (2022) look at AI adoption by specifically focusing on firm investments in data training. Start-ups invest in developing new algorithms, among which are neural networks and ensemble learning, and the development of AI-enabled novel products. To do so, they need incentives to invest in addition to mitigating their financial barriers. In line with the core of most of the innovation literature, Bessen et al. (2023) find that US startups invest in data training and are able to attract VC capital only when they can count on proprietary data, rather than publicly available data. While this result is not particularly surprising, it does raise important issues from the policy perspective. First, investment in data and data training is more likely to lead to market concentration, calling for a rethinking of competition policy in the context of digitalization (Furman et al., 2019). Second, not all types of data are intrinsically subject to firm property rights, particularly when the collection and use of personal data are involved (Savona, 2020), as we will see in Section III.

Also in this special issue, Igna and Venturini (2023) focus on the impact of AI adoption on firm performance across a large sample of European firms over the period 1995–2016. Drawing on patent data, and in line with most of the innovation literature, they show that the inventive success of firms in the AI realm strongly relies on path dependency, that is, on prior investments in Information and Communication Technologies. They find high complementarity and learning spillover effects between investments in AI and the accumulated knowledge in data management, network and communication and, most recently, cognition and imaging.

3. The impact of AI on work

This section focusses on the opportunities and challenges from AI – and other emerging digital technologies – for labour markets. AI has the potential to increase productivity, employment, and wages. However, while in the four decades immediately after WWII technological progress seemed to result in better outcomes for both high-paid as well as low-paid workers, a very different impact of technological progress

² The project has harmonized official micro survey data across 11 countries (Belgium, Denmark, France, Germany, Ireland, Israel, Italy, Japan, Korea, Portugal and Switzerland), which go beyond the US or EU only focus.

started to emerge in the 1980s, which was much less inclusive for low-paid work.

Designing policies to tackle this challenge requires a proper understanding of how technological progress has impacted on labour markets in the recent past, and how AI is likely to change our jobs in the future. Autor (2022) provides an overview of the recent opinions of the impact of digital technologies on labour markets. His starting point is the “task-based view of labour markets”, which has become the standard framework in the recent literature.

The hypothesis put forward by this task-based view is that digital technologies can automate “routine tasks”. What makes a task routine is that it follows an explicit, fully specified set of rules and procedures. Consequently, routine tasks can in many cases be codified in computer software and executed by machines (e.g., robots to assemble a car, email to deliver messages). Conversely, “non-routine tasks” have historically been challenging to program because the explicit steps for accomplishing these tasks are often not formally known or they imply the use of tacit knowledge (Cowan et al., 2000; Foray and Steinmueller, 2003). Paradoxically, even though we cannot formally express non-routine tasks in an algorithm, many of these tasks are easy for humans to do. This is known as Polanyi’s paradox (“humans know more than they can tell”), named after 20th century philosopher Michael Polanyi and his argument that all our knowing is rooted in tacit knowledge.

Goos et al. (2014) show that routine tasks are concentrated in middle-paid occupations (e.g., machine operators, office clerks), while non-routine tasks are concentrated in low-paid occupations (e.g., waiting tables in a restaurant, cleaning a room) and high-paid occupations (e.g., medical diagnosis, managing a team). Consequently, digital technologies have decreased demand for middle-paid occupations, relative to low-paid and high-paid ones, resulting in a process of job polarization. They show that this is happening in the 16 (Western) European countries that they examined from 1993 to 2010, and similar evidence exists for the US (Acemoglu and Autor, 2011).

Job polarization is linked to changes in relative wages and inequality. Acemoglu and Restrepo (2022) show that between 50 % and 70 % of changes in the US wage structure over the last four decades are accounted for by relative wage declines of worker groups specialized in routine tasks in industries experiencing automation. Emerging literature also examines the impact of firm-level investments in automation rather than actual exposure or adoption of digital automation technologies - with the only exception of Prytkova et al. (2023) - on firm-level outcomes such as its labour share or occupational structure. In this light, based on French administrative and matched employer-employee data, Domini et al. (2022) in this special issue examine the impact of investments in, and importantly also import of, AI-related goods on within-firm wage inequality. They find little evidence of spikes of investment and import of AI-related goods having any impact on relative wages and gender-wage inequality within firms. The authors detect an effect that is driven mainly by new hirings rather than on incumbent workers, with little direct effect on gender wage-gap within firms.

Looking ahead, AI is likely to fundamentally change the relationship between digital technologies, labour demand, and inequality. AI can be used to infer tacit relationships that need not be fully specified by underlying software because AI learns to perform tasks inductively by training on examples instead of by following explicit rules that are programmable. Basically, in the present AI-era, technology can readily acquire the tacit knowledge that programmers in the pre-AI era struggled to embed in computers.

Consequently, many non-routine tasks done in both low-paid and high-paid occupations could be performed by AI in the future, resulting in very different changes in labour demand and inequality than we have seen over the past decades. For example, we might no longer see a process of job polarization but of stronger relative employment growth in high-paid occupations (if AI automates non-routine tasks in low-paid but not high-paid occupations) or of stronger relative employment growth in low-paid occupations (if AI automates non-routine tasks in

high-paid but not low-paid occupations).

Because of AI’s promise of a paradigm shift in our thinking about its impact on employment, wages and inequality, there is much uncertainty around AI’s implications for labour markets. The remainder of this section focusses on the following four questions:

1. Which worker tasks will AI automate?
2. Which new worker tasks will emerge from AI?
3. What will be the impact of AI on workers?
4. What will be the impact of AI on working conditions?

3.1. Which worker tasks will AI automate?

While earlier digital technologies automated occupations intensive in doing routine tasks (e.g., machine operators, office clerks), machine learning as a prediction technology has the potential of also automating various non-routine tasks across a wide range of occupations. To study this question, small but rapidly growing literature has emerged that applies a task approach to analyze the impacts of AI adoption on different occupations (Acemoglu et al., 2022b, Brynjolfsson et al., 2019, Felten et al., 2019, Webb, 2020). Although these studies use the task framework generally, they do not start from the premise that AI can only do a given set of tasks. Instead, they rely on various innovative ways to determine which tasks AI can(not) automate.

Webb (2020) offers one example. He uses Natural Language Processing (NLP) algorithms that exploit the overlap between the text of job task descriptions and the text of patents to develop a new method for identifying which tasks can be automated by any technology. This allows him to construct a measure of the “exposure” of occupations to that technology. For example, a doctor’s job description includes the task “diagnose a patient’s condition”. An NLP algorithm then extracts the verb-noun pairs from this task, being “diagnose condition”. The algorithm then quantifies the same verb-noun pairs in a different corpus of patents to identify whether any technology could automate a doctor’s job.

Using this approach, Webb (2020) first examines the impact of two previous types of digital technologies: traditional software and robots. For software, exposure is decreasing with education, with individuals in middle-wage occupations being most exposed. Men are much more exposed to software than women, reflecting the fact that women have historically clustered in occupations requiring complex interpersonal tasks which software has not been capable of doing. For robots, individuals with less than a high-school education and men under age 30 are most exposed, consistent with robots’ substituting for middle-paid manufacturing jobs. By and large, these results are consistent with the literature on job polarization, which has found that digital technologies reduced demand for routine middle-wage jobs while increasing it for non-routine low- and high-wage jobs between 1980 and 2010.

Webb (2020) then turns to the impact of AI on the demand for occupations. In contrast to traditional software and robots, AI performs tasks that involve detecting patterns, making judgments, and optimization. Most-exposed occupations include clinical laboratory technicians, chemical engineers, optometrists, and power plant operators. More generally, high-skilled occupations are most exposed to AI. There is also a small proportion of low-skilled jobs that are highly exposed to AI. Examples are production jobs that involve inspection and quality control. In sum, the impact of AI on labour markets is likely to be very different from previous advances in digital technologies. This suggests a paradigm shift in our thinking about AI’s potential to automate worker tasks.

Savona et al. (2022) look at a large pool of technical and engineering papers describing the ideation, design and prototype development of digital automation technologies, specifically Robots, AI, Data Management and Data Acquisition. They uncover how these technologies are designed for different sectors and tasks and summarise the narratives on

how human-machine interactions develop to substitute or complement humans across routinized and non-routinized tasks. Emerging digital automation technologies, even within the same family, are intrinsically heterogeneous in their design and the tasks they can execute, though the number of sectors that are exposed to them is still relatively limited. Data-intensive technologies are more pervasive in services and complementary to human activities, while robots and AI prevail in the manufacturing sectors, with comparatively less interaction with humans.

3.2. Which new worker tasks will emerge from AI?

If the set of tasks were static, automation by AI would crowd workers into an ever-narrowing subset of tasks, perhaps finally making human labour altogether obsolete if AI would evolve into a state of AGI. However, it is more likely that even AGI will create many new jobs for workers.

Nevertheless, very little is known about how many new jobs any type of technological progress (including AI) creates. To answer this question, [Autor et al. \(2022\)](#) exploit the emergence of new job titles in the US Census Bureau's occupational descriptions that survey respondents in their Census forms. Their analyses show that, irrespective of whether a new job is created because of technological progress or some other reason, new work is quantitatively important. They estimate that more than 60 % of US employment in 2018 was in occupations that did not exist in 1940. Examples of new occupational titles are "fingernail technician", which was added in 2000, and "solar photovoltaic electrician", which was added in 2018.

Regarding the nature of new work, they find that between 1940 and 1980 most new work that employed non-college workers was found in middle-skilled occupations. After 1980, however, the locus of new work creation for non-college workers shifted away from these middle-tier occupations and moved toward traditionally lower-paid personal services. Conversely, employing new, college-educated workers became increasingly concentrated in professional, technical, and managerial occupations. In combination, these patterns indicate that new work creation polarized after 1980, mirroring (and in part driving) aggregate job polarization.

To further explain the origins of new job titles, and the role of technological progress, [Autor et al. \(2022\)](#) follow a procedure like [Webb \(2020\)](#) by examining patent data using NLP. Unlike [Webb \(2020\)](#), however, they instructed their NLP algorithm to look for text that indicates augmentation instead of automation of worker tasks. For example, in 1998, the U.S. Patent and Trademark Office (USPTO) granted a patent for a "method of strengthening and repairing fingernails". Their algorithm links this patent to the occupational title "fingernail technician", which the Census Bureau added in 2000. Similarly, their algorithm links the 2014 patent "systems for highly efficient solar power conversion" to the occupational title of "solar photovoltaic electrician", which was added in 2018. In sum, [Autor et al. \(2022\)](#) show that new technologies are an important driver for the creation of new work.

[Autor et al. \(2022\)](#) also instruct their NLP algorithm to look for text in patents that indicate a new technology's potential to automate (instead of augment) worker tasks. For example, the introduction of photocopying in part automated the jobs of office clerks. Moreover, they find that technology has not been able to augment the job of office clerks in the same way, leading to a net decrease in labour demand and employment for this occupation. Conversely, in other occupations, such as mechanical engineers or operations research analysts, augmentation has been more important than automation, resulting in an increase of employment in these occupations. Interestingly, they also show that occupations are either simultaneously exposed to both augmentation as well as automation, or not exposed to any technology at all.

[Prytkova et al. \(2023\)](#) offer a novel methodological framework based on NLP to estimate industrial and occupational exposure scores to

emerging Science, Technology and Innovation area, distinguishing between user and producer industries. They rely on over 190,000 novel patents mapped into clusters of over 500 emerging STIs, 271 (3-digit NACE) industries and 433 (4-digit) ISCO-08 tasks in 28 European countries from 2011 to 2019. They find that occupations that are more (less) exposed to emerging STIs are associated with increased (decreased) levels of employment between 2012 and 2019, with occupations mainly involving non-routine tasks being the top ones exposed. Interestingly, high occupational exposure in technology-producing sectors is associated with employment growth, while low exposure is associated with employment decline. Technology-user sectors show an opposite trend: the cluster of deep IT users reveals a decline in employment associated with exposure.

In conclusion, while technology automates jobs, it also augments work and is an important driver of new job creation. [Autor \(2022\)](#) coins this double-sided impact of innovation on work "the race between automation and augmentation". In occupations with declining (increasing) employment shares, this race is won by automation (augmentation). Better understanding this race gives policymakers important levers to seize the benefits from AI. For example, a race between automation and augmentation of worker tasks implies that AI has the potential and can perhaps be steered toward more augmentation and less automation.

The impact of technological progress, including AI, on work is characterized by competing forces of automation and augmentation of worker tasks. This is true even within narrowly defined occupations. The focus of researchers, as well as managers, entrepreneurs, and policymakers, should therefore be not (just) on AI's automation or augmentation potential, but also on job redesign. For example, [Brynjolfsson et al. \(2018\)](#) conjecture that machine learning will require a substantial redesign of tasks for concierges, credit authorizers, or brokerage clerks.

In this special issue, [Belloc et al. \(2022\)](#) tackle the specific challenge of how to direct the governance of digital automation within workplaces to ensure an outcome of augmentation rather than replacement through automation. Interestingly, based on an evolutionary model that represents workplace governance with different degrees of Employee Representations (ER), they find that automation risk is higher when the governance of workplaces does not allow for employee representation. When, instead, workplaces incorporate ER, it is possible to steer job design toward higher augmentation and lower automation. Their paper also offers an empirical test of the model, based on a large sample of European workers, and concludes that it is crucial for labour market policies to be sensitive to the socio-institutional aspects that might steer the directions of technical change toward positive synergies between humans and machines. Involving workers in the co-design of jobs is determinant. More generally, it is argued that the future of work as a result of the adoption of digital automation technologies should be understood not only in terms of technological adoption, and task redefinition, but also in terms of organizational choices and institutional settings, which can be more or less ER friendly. Arguably, as pointed out by [Belloc et al. \(2022\)](#), the effectiveness of ER in the context of a fast-changing technological landscape might be reduced by the emergence of alternative work arrangements and self-employment ([Ciarli et al., 2018](#); [Ciarli et al., 2020](#)).

This line of research can be further developed both theoretically and empirically to understand the role that industrial relations can play at the firm and national level to steer the impact of AI on labour toward augmentation rather than job displacement.

3.3. What will be the impact of AI on workers?

3.3.1. Worker skills

The need for job redesign highlights the need for on-the-job re-learning by workers. Using the new module introduced in the US Census Bureau's 2019 Annual Business Survey (ABS), which we discussed in the previous section, [Acemoglu et al. \(2022\)](#) examine firms' self-assessment

of the effects of AI on their demand for workers' skills. Half of AI adopters report that advanced technologies increase their skill demands, while almost no firms report a reduction. This self-reported increase in firms' skill requirements when they adopt AI explains part of the well-known skills gap and highlights the importance of on-the-job investments in worker skills.

Genz et al. (2021) provide similar evidence for Germany. They examine how German workers adjust to firms' investments in new digital technologies, including AI, augmented reality, or 3D printing. To do this they collected novel data that links survey information on firms' technology adoption to administrative social security data for Germany. They then compare technology adopters relative to non-adopters. While they find little evidence that AI affected the number of jobs, the absence of an overall employment effect masks substantial heterogeneity across workers. They find that workers with vocational training receive more benefit than workers with a college degree. One explanation might be that AI augments vocational work more than it augments tasks done by college workers. Another explanation is that Germany's traditionally strong vocational training system (76 % of all workers completed vocational education) provides an abundance of specialized skills that directs the development and adoption of AI toward making use of (and thereby augmenting) vocational skills.

3.3.2. Worker mobility

Some workers will be displaced from their jobs because AI automates much more than it augments the tasks they do. Job displacement is costly for those made redundant, and it could be disruptive for the labour market more generally. These adjustment costs and disruptions are characteristic of previous episodes of technological change. And given the speed with which AI is evolving, they may become particularly acute. However, very little is known about the transition of workers who are displaced from their jobs.

One exception is Bessen et al. (2023). Using Dutch administrative data, they examine what happens to workers made redundant when their firm invests in the automation of its existing procedures. They find that the expected annual income loss across all workers before their firm adopts automation technologies accumulates to 9 % of one year's earnings after 5 years. They also show that this annual income loss is driven by spells of unemployment within a year (rather than, for example, quickly moving into lower paid jobs), with unemployment benefits only insuring partially against their income losses. These adverse impacts of automation are larger in smaller firms and for older and middle-educated workers. In sum, their results suggest that there are substantial adjustment costs for displaced workers, and that these adjustment costs are only partially offset by unemployment insurance.

Another exception is Feigenbaum and Gross (2022), who examine the introduction of mechanical switching in operating telephone calls that took place in half of all US states between 1920 and 1940. They study adjustments in the labour market for young female telephone operators, one of women's main occupations at the time. They find that telephone operators were significantly less likely to still be working as operators ten years after their state's cutover to mechanical switching. While some found other jobs in the telephone industry, others (especially older workers) left the workforce, and those who remained employed were more likely to have switched to lower-paying occupations. They also find that automation of telephone operating did not decrease overall demand for young women in their local labour markets. After automation, young women were less likely to become telephone operators but entered different jobs such as middle-skilled clerical and lower-skilled service occupations (mainly typists and waitresses).

3.3.3. Algorithmic hiring of workers

During the past decade, many firms transformed their hiring practices. While the central goal of hiring remains the same, the set of tools available has changed, primarily due to innovations in AI. AI can screen résumés on a massive scale to discard applicants that are likely to be a

poor fit; it can then put the remaining candidates through assessments to further narrow the list of suitable candidates. For many firms, only at the later stages of the process do humans enter the picture: final interviews, negotiations, and convincing a candidate to accept an offer remain important tasks for HR professionals.

Li et al. (2020) examine biases in algorithmic hiring. They find that firms hiring workers must balance "exploitation" (selecting job applicants from groups with proven track records) with "exploration" (selecting job applicants from under-represented groups to learn about quality). They show that a learning algorithm that screens résumés and values both exploitation but also exploration improves the quality of candidates selected for an interview, while also increasing demographic diversity. Their results highlight the importance of incorporating AI in hiring processes to make them both more efficient and equitable.

In this special issue Fumagalli et al. (2022) use a novel approach to examine biases in algorithmic recruitment and the effects on labour markets and workers' welfare. They focus on the perception of workers rather than that of employers by using algorithmic recruitment to save costs or improve accuracy in the evaluation of workers' task performance. The authors carry out two incentivized experiments to show whether workers prefer algorithmic or human assessment in the recruiting process. Interestingly, human recruiters are found to be more error-prone and less consistent in task performance evaluation, while putting more weight on personal characteristics. Hence, high performers prefer algorithmic assessment, while lower task performers prefer human evaluation. Another intriguing finding is that preferences for human assessment are dictated by suggestive evidence of gender bias in evaluations performed by machines versus humans.

Fumagalli et al. (2022) add to several ongoing debates, which go beyond potential biases in assessment, and touch the more general issue of increasing machine-led surveillance of workers' performance, particularly platform workers (Baiocco et al., 2022) and the personal data regulation issues that this entails (Adams-Prassl, 2019).

A further line of research that we hope this work can inspire is the extent to which algorithmic recruitment can (or cannot) properly evaluate 'soft skills' in a context where these seem to become more important, particularly in public and social services and for workers at the bottom of the skills and wage distribution (Aghion et al., 2022).

In summary, arguably the issue of biases in recruitment and the workers' perception is a crucial one, that needs careful address in the future, both at the workplace level and the national institutional level.

3.4. What will be the impact of AI on working conditions?

Wood (2021) discusses the prevalence of algorithmic management of workplaces. Algorithmic management relies on data collection and surveillance of workers to manage workforces in an automated way. A well-known example is online labour platforms (see Fernandez Macias et al., 2023 for a comprehensive review). These platforms enable workers to choose the clients and jobs they take, how they carry out these jobs, and the rates they charge to do them. However, to varying degrees, workers' ability to make these choices is strongly shaped by platform rules and design features. Increasingly, algorithmic management is also being used in other settings, such as in warehouses, retail, manufacturing, marketing, consultancy, banking, hotels, call centres, and among journalists, lawyers, and the police. Wood (2021) summarizes several detailed case studies from these sectors.

Increasing attention has been devoted to a particular segment of labour markets, that of platform and online workers, or gig workers, as they are usually referred to (Fernandez-Macias et al., 2023). This attention is due to these types of workers being particularly vulnerable to phenomena such as algorithmic surveillance, as mentioned earlier, precarity of their non-standard work arrangements (Ciarli et al., 2020), punitive wage conditions, and lack of representation and protection, among other characteristics. Platform workers are service providers that deal directly with platform owners, which in turn enjoy significant

network economies and bear negligible transaction costs (Mansell and Steinmueller, 2020).

In this special issue, Duch-Brown et al. (2022) focus on Online Labour Markets (OLM), which is the ensemble of contractual terms and conditions designed and applied by employers on large platforms to their (online) employees. They look at the roots of OLM's market power to rule, design features of workers' tasks, their wages, and more generally, the characteristics of labour supply and demand elasticities. Their empirical analysis is based on matched data on about 500,000 digital projects, about 200,000 hiring employers (labour demand) using the PeoplePerHour (PPH) platform across 180 countries and applications by over 100,000 workers (labour supply) across the same countries. Their estimation on labour demand and supply elasticities, in addition to market shares, confirm the strong and dimensionally large market power of the platform and its effect on wage settings of workers. However, when distinguishing between demand and supply respectively of AI and non-AI workers, they find a relatively scarce AI-skilled labour force, which enjoys a wage premium of over 3 % compared to the non-AI skilled digital workforce.

Weil (2014) discusses the broader impact that algorithmic management has on business models and labour relations. He argues that firms can use information and communication technologies to erode the need for traditional employment relations. Since the 1980s, "lead companies" have shed their role as direct employers in favour of outsourcing work to smaller subcontractors or franchisees. Competition between these subcontractors or franchisees implies that costs, including wages, are lower compared to when lead companies would directly employ these outsourced workers. Weil (2014) refers to this as "the fissuring of workplaces". Because fissuring of workplaces mainly affects low-wage jobs, it has contributed to higher wage inequality and to increased occupational safety and health risks for workers in fissured jobs. A key enabler of fissured workplaces has been advances in information and communication technologies, and recent developments in AI could enable lead companies and their shareholders to manage their labour supply chains even better through intelligent monitoring of outsourced workers.

The issue of employment relations in the context of platform owners and workers, particularly in terms of the consequences of algorithmic management on workers and the opportunities to regulate it, is explored in this special issue by Rolf et al. (2022). In the context of the UK regulatory framework of Social and Employment Protection (SEPs) and based on the case of the UK courier sector, the authors draw on socio-legal theories to unpack the challenges that platform workers face in the presence of regulatory voids. The case of the 'self-employment plus' is put forward as an example of how privatized social protection systems are being developed to fill such a void, with consequences of (further) weakening of platform workers protection. In the absence of a proper national and public regulatory framework, the authors argue, it is likely that the platform ecosystem sees regulatory oversight, diminished social protection, increases in the algorithmic managerial control and surveillance and productivity extractives practices exacerbated.

The regulatory and governance implications of these findings can be summarized as follows. First, it is imperative to fill the regulatory void that still exists concerning platform workers and online labour markets. Second, worker representation and a stronger voice in the co-designing of digital projects and tasks, as well as bargaining power on remunerations, should be facilitated. The authors of this article put forward several specific policy recommendations and the need to equalize standard and non-standard work arrangement already a few years ago in the context of the EC High Level Expert Group on the impact of digitalization on European Labour Markets (European Commission, 2019). Some of these recommendations, such as the Digital Single Window for platform workers, have been explored further by European policymakers, while others remain to be implemented.

4. Data governance

Economists of innovation have included investments in database among the innovation expenditures since the very first editions of the Oslo Manual (OECD, 1992). Investments in data collection, aggregation, analytics and data management skills are included among the intangible assets of firms (Corrado et al., 2009). Despite their large diversity, investments in intangibles are argued to contribute to the knowledge-based capital in firms, including Research & Development (R&D), Intellectual Property Rights (IPR), training, software, engineering and design, marketing and branding. More recently, intangibles and database have been increasingly considered in national accounts (OECD, 2019b; Ahmad and van de Ven, 2018).

Arguably, though, a theoretical grounding of the economic nature and the value of data, starting from personal data, is still lacking, despite some attempts (Acemoglu et al., 2022c)

Personal data³ is a *club good*, excludable but not rivalrous (Savona, 2019), as individuals might prevent the sharing and use of their personal information; though once released, data can be "used multiple times (e. g. in different contexts) without inherently diminishing their value. In principle, data can be exploited and re-exploited infinitely at low marginal cost; it is data infrastructure and analytics that are the primary costs related to data re-use" (OECD, 2019a, p. 240).

The economic nature of data changes along the data 'value chain'. Firms invest in digital infrastructure and analytics, they collect, store and analyze databases and are granted property rights on it.⁴ Non-rivalry of shared data allows firms investing in database to benefit from economies of scale and network economies, as in the case of large platforms (Mansell and Steinmueller, 2020). While individual data is a club good, a database is a *private good*, excludable and rivalrous, and provides information that represents a comparative advantage for firms. Data analytics ensuing from aggregate information becomes knowledge, that is inherently a *public good* (Foray, 2004). Non-rivalry and network economies also allow public and non-profit institutions to harness social value from data, when the knowledge stemming from aggregating information serves a public purpose.

The literature on digital economics (see among others Goldfarb and Tucker, 2019) has looked at issues of privacy, exploring the boundaries of property rights to be attributed to data, from a firm and an antitrust perspective, considering data as an 'essential facility' (Tucker, 2019).⁵

The changing economic nature of data along the data value chain makes the attribution of rights and the governance of data particularly challenging. From a policy perspective, one challenge is to reconcile objectives that are often at odds with each other: creating incentives for

³ Personal data means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person" (article 4(1), EU GDPR, 2018).

⁴ EU Directive 96/9/EC of 11 March 1996 recognizes the legal ownership of database to firms, with *database property rights* being a legal category implemented in that context.

⁵ Westin (1968) had already defined privacy as "the claim of individuals, groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others" but also that "personal information (...) should be defined as a property right, with all the restraints on interference by public or private authorities (...) that our law of property has been so skillful in devising". Nowadays, data protection laws focus on privacy safeguards; they do not imply active use rights (Specht and Zerbst, 2018). In Europe, privacy is a fundamental right, equated to human dignity; it is inalienable and cannot be sold or traded (van Lieshout, 2015; Lynskey, 2015; Zander et al., 2019). However, firms can and have long been able to monetize personal data as they are granted property rights over database.

maximising data sharing for purposes of public interests, such as health, green mobility or research; *capping private value concentration* ensuing from (unintended or voluntary) data sharing, for instance to large platforms, while *protecting individual privacy and other rights* (Savona, 2019, 2020). The European Commission has been at the forefront to resolve this policy conundrum in the context of the articulated regulatory framework developed over the past few years, concerning B2B and B2G data sharing. The European regulatory framework on data is considered a benchmark worldwide, to the extent that several countries such as India, Canada, and some of the states in the US are adopting General Data Protection and AI regulations based on the EU model. Many steps have in fact been taken at the European level, including, among others, the recent Data Act (European Commission, 2023). One of the most distinctive features of the EU regulatory approach to data governance is the creation of incentives for data sharing.

For instance, let us consider the EU Data Governance Act (DGA) (European Commission, 2022). The DGA explicitly aims to foster the “availability of data for use by *increasing trust in data intermediaries* and by *strengthening data sharing mechanisms* across the EU.” The main focus is in the *creation of data markets* by legitimising data intermediaries (i.e. data trusts, cooperatives, stewards, unions). Further, it aims to “*make public sector data available for re-use (...) on altruistic ground.*”

When considering the economic nature of data, though, Some questions arise (Savona, 2020), that might have implications on the effectiveness of the regulation:

- What is the minimum scale for data intermediaries that manages individual data for altruistic purposes to be effective? Altruism and public interest, such as research, public health, and some public services, need a large-scale data sharing infrastructure (at least a minimum benchmark) to be effective. On the other hand, what is the maximum scale that we are to allow data intermediary services to avoid risks of value concentration, privacy leakages, and cyber security, which we all know too well from the experience with large tech?
- What ensures that data intermediaries or trustees that operate on a fiduciary basis on behalf of a group of individuals would have incentives compatible with altruism? If this is not the case, then the risk is to shift from big tech to big trusts. These need to be necessarily capped in scale, limited to specific purposes and monitored by an independent, governing body that includes representatives from the relevant data subject constituencies (or relevant sector-level data) rather than a proliferation of small bottom-up data intermediaries in a fragmented market.

Mahieu (2021) argues for the right to *informational self-determination* while trying to reconcile the right to privacy and that of property. A thorough reflection on the economic nature of data and the appropriate boundaries of property and privacy rights to tackle the asymmetries in data markets (i.e. due to different levels of digital literacy, access to information, regulatory and institutional capacity) is needed. Would a proliferation of data intermediaries ensure this? Or, rather, would a regulated public actor, monitored by an independent guarantor and representative of all the different constituencies and stakeholders in societies, achieve the objective of a just data governance?

In this special issue, Graef and Prüfer (2021) propose a governance framework that allows data sharing, in the case of a B2B relationship, which is data sharing across data intensive firms, for the purpose of avoiding market tipping and market concentration. From a legal perspective, they claim that data sharing should be made mandatory and regulated through three potential models. The first model would be a fully centralized one, involving the creation of a European Data Sharing Agency that would enforce data sharing centrally; the second model would be fully decentralized, involving the setting of a Data Sharing Cooperation Board that oversees a network of National Competition Authorities (NCAs) that are demanded to enforce data sharing contracts;

the third one would be a mixed model, with centralized and decentralized features of delegating investigation and decision making to national authorities and with an overseeing (EU) cross-national Data Sharing Agent that is competent in enforcement.

However, as most data governance issues arise because they impinge upon the use of personal data or processed data that aggregates and anonymises personal and potentially sensitive information, it would be advisable that all relevant stakeholders are included in the decisions and possibly the design of data governance instruments. These include not only major private players, but also local governments and civil society. Making data available for the public interest through mandatory data sharing or creating incentives for data sharing that do not lower consumer and citizens' protection is not an easy task. Arguably, data governance is one of the realms where the aid of citizen science (Beck et al., 2022) is particularly valuable. Participatory methods such as ‘Citizen Juries’ are valuable instruments to inform and involve the public in understanding the main trade-offs of policy decisions concerning a just data governance.

Gravey et al. (2023), for instance, look at the result of a sample of citizens deliberating and voting over a simulated trade policy intervention that concerned medical data sharing. Citizens considered that lowering trade barriers to cross-border data flows and facilitating international sharing of health and medical records might allow improvements in digital health services, advances in medical treatments in areas that require the collection and analysis of high volumes of data, such as rare diseases, and facilitate innovations in health e.g. by pharmaceutical companies. However, such a decision might lead to sharing health data with countries with lower privacy and data protection standards than the UK, potentially resulting in privacy abuses, negative effects on those most at risk from this data being shared (such as those with long-term health conditions and refugees), the commercial profiling of consumers more generally. The deliberations have shown that citizens significantly rely on “independent experts”, but they also attached value to their personal records and were prone to share only when this entailed high returns in terms of social value.

Arguably, a broader conceptualization, design and implementation of data governance, in a context where the use of AI is spreading at an unprecedented speed, is required and needs multidisciplinary expertise. This special issue hosts technological, economic and legal contributions to the governance of data and AI. The case of generative AI mentioned in the introduction is an example of the extent to which it is important to understand and predict the concerns that emerging digital automation technologies raise, some of these have been unprecedented in the history of other technological paradigms. For instance, never have the very entrepreneurs and innovators, owners of ‘too big to fail’ platforms, emphatically demanded regulatory intervention from governments to ‘slow down’ the development of generative AI⁶. Neither have they explicitly expected a public body to identify and regulate undesirable effects such as fake news, misinformation and cyber security. It has never happened before that intellectual property protection, including copyright, was allegedly violated by a machine, or - on the contrary - attributed to an algorithm. This is an area of research that would deserve a substantial multidisciplinary effort to resolve the current pressing, and potentially increasing, challenges around data governance.

5. Conclusions

The design and implementation of digital industrial policy, particularly in a phase of pervasive digital transformation, requires the understanding of how digitisation change the delivery of public and private services and directly affect the organization and structure of product and labour markets. Regarding labour markets, it is not only a matter of

⁶ See “Pause Giant AI Experiments: An Open Letter (March 2023): <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>.”

investing in new skills and anticipating the kind of occupations and tasks that digital technologies require; it is also important to understand how digitisation changes the very notion of creativity and the organization of work, while generating new power asymmetries and modifying working conditions, workers' representation and the protection of acquired rights.

In the context of 'sustainable' digital governance, ensuring an inclusive governance of data becomes increasingly important. This will require redefining the boundaries of knowledge as a public good, the sustainable creation of incentives for data sharing and support to the development of collective intelligence in addition to the simple aggregation of information. Data is key and its economic nature makes the allocation of rights to the actors involved particularly challenging. One of the challenges of data governance is to reconcile often conflicting objectives: to create (and maintain) incentives to maximise data sharing for purposes of public interest, such as health or research; to limit the concentration of private value arising from (involuntary or voluntary) data collection and analytics; to protect individual privacy and other rights such as copyright in a context where human creativity (still) has a social value.

Europe will also have to face the challenges of how to govern international digital trade. In this context, on the one hand, the structure of European regulations focusing on the protection of individual rights, and representing a global benchmark, must be preserved. On the other hand, the flow of international data transfers must be regulated to preserve EU fundamentals in a context of global geopolitical tensions.

These reflections should be embedded in the context of a global AI regulation and digital governance. The EU, the US and China are adopting diverging regulatory frameworks, and the resulting regulatory fragmentation and asymmetries in data and AI governance, particularly when dealing with cross-border data flows or data sharing with different levels of data protection, should not offset the main objective of using AI to support a just digital transition.

Acknowledgements

Maarten Goos thanks the Instituut Gak for financial support. Maria Savona gratefully acknowledges the European Union's Horizon 2020 Research and Innovation Programme under Grant agreement ID 101004703 PILLARS (Pathways to Inclusive Labour Markets) and the ESRC Centre for Inclusive Trade Policy at the University of Sussex.

References

- Acemoglu, D., 2021. Harms of AI. In: NBER Working Paper 29247. National Bureau of Economic Research, Cambridge, MA. <https://www.nber.org/papers/w29247>.
- Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: implications for employment and earnings. In: *Handbook of Labor Economics*, 4, pp. 1043–1171. [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5).
- Acemoglu, D., Restrepo, P., 2022. "Tasks, Automation and the Rise in US Wage Inequality" Forthcoming in *Econometrica*. <http://pascual.scripts.mit.edu/research/taskdisplacement/>.
- Acemoglu, D., Autor, D., Hazell, J., Restrepo, P., 2022a. Artificial intelligence and jobs: evidence from online vacancies. *J. Labor Econ.* 40 <https://doi.org/10.1086/718327>.
- Acemoglu, D., Anderson, G., Beede, D., et al., 2022b. Automation and the workforce: a firm-level view from the 2019 annual business survey. In: Paper Presented at the NBER/CRIW Conference on Technology, Productivity and Economic Growth, Washington DC, March 2022. <http://pascual.scripts.mit.edu/research/abs/>.
- Acemoglu, D., Makhdoumi, A., Malekian, A., Ozdaglar, A., 2022c. Too much data: prices and inefficiencies in data markets. *Am. Econ. J.* 14 (4), 218–256. *Microeconomics*.
- Adams-Prassl, J., 2019. What if your boss was an algorithm? The rise of artificial intelligence at work. *Comp. Labor Law Policy J.* 41 (1), 123. Available at SSRN: <https://ssrn.com/abstract=3661151>.
- Aghion, P., Bergeaud, A., Blundell, R., Griffith, R., 2022. *Soft Skills and the Wage Progression of Low-educated Workers*. Mimeo University College London.
- Ahmad, N., van de Ven, P., 2018. *Recording and Measuring Data in the System of National Accounts*. OECD, Paris.
- Autor, D., 2022. The labor market impact of technological change: from unbridled enthusiasm to qualified optimism to vast uncertainty. In: NBER Working Paper 30074. National Bureau of Economic Research, Cambridge, MA. <https://www.nber.org/papers/w30074>.

- Autor, D., Chin, C., Salomons, A., Seegmiller, B., 2022. New frontiers: the origin and content of new work 1940–2018. In: NBER Working Paper 30389. National Bureau of Economic Research, Cambridge, MA. <https://www.nber.org/papers/w30389>.
- Baiocco, S., Fernández-Macias, E., Rani, U., Pesole, A., 2022. *The Algorithmic Management of Work and Its Implications in Different Contexts*. European Commission, Seville, p. JRC129749.
- Beck, S., Brasseur, T.-M., Poetz, M., Sauermaun, H., 2022. Crowdsourcing research questions in science. *Res. Policy* 51 (4). <https://doi.org/10.1016/j.respol.2022.104491>.
- Belloc, F., Burdin, G., Cattani, L., Ellis, W., Landini, F., 2022. Coevolution of job automation risk and workplace governance. *Res. Policy*. ISSN: 0048-7333 51 (3), 104441. <https://doi.org/10.1016/j.respol.2021.104441>.
- Bessen, J., Michael Impin, S., Reichensperger, L., Seamans, R., 2022. The role of data for AI startup growth. *Res. Policy*. ISSN: 0048-7333 51 (5), 104513. <https://doi.org/10.1016/j.respol.2022.104513>.
- Bessen, J., Goos, M., Salomons, A., Van den Berge, W., 2023. "What Happens to Workers at Firms That Automate?" Forthcoming in *Review of Economics and Statistics*. https://doi.org/10.1162/rest_a_01284/114750/What-Happens-to-Workers-at-Firms-that-Automate?redirectedFrom=fulltext.
- Bianchini, S., Müller, M., Pelletier, P., 2022. Artificial intelligence in science: An emerging general method of invention. *Res. Policy*. ISSN: 0048-7333 51 (10), 104604. <https://doi.org/10.1016/j.respol.2022.104604>.
- Brynjolfsson, E., Mitchell, T., Rock, D., 2018. What can machines learn, and what does it mean for occupations and the economy? *AEA Pap. Proc.* 108, 43–47. <https://doi.org/10.1257/pandp.20181019>.
- Brynjolfsson, E., Collis, A., Diewert, E., Eggers, F., Fox, K., 2019. GDP-B: accounting for the value of new and free goods in the digital economy. In: NBER Working Paper 25695. National Bureau of Economic Research, Cambridge, MA. <https://www.nber.org/papers/w25695>.
- Calvino, F., Fontaneli, L., 2023. A portrait of AI adopters across countries: firm characteristics, assets' complementarities and productivity. In: *OECD Science, Technology and Industry Working Papers*, No. 2023/02. OECD Publishing, Paris. <https://doi.org/10.1787/0fb79bb9-en>.
- Ciarli, T., Marzucchi, A., Salgado, E., Savona, M., 2018. The Effect of R&D Growth on Employment and Self-employment in Local Labour Markets (March 23, 2018). SWPS 2018-08. Available at SSRN: <https://ssrn.com/abstract=3147861>.
- Ciarli, T., Di Ubaldo, M., Savona, M., 2020. Innovation and self-employment. In: Zimmermann, K. (Ed.), *Handbook of Labor, Human Resources and Population Economics*. Springer, Cham. https://doi.org/10.1007/978-3-319-57365-6_17-1.
- Cockburn, I., Henderson, R., Stern, S., 2019. The impact of artificial intelligence on innovation: an exploratory analysis. In: Agrawal, A., Gans, J., Goldfarb, A. (Eds.), *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, Chicago, pp. 115–146. <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda/impact-artificial-intelligence-innovation-exploratory-analysis>.
- Corrado, C., Hulten, C., Sichel, D., 2009. Intangible capital and US economic growth. *Rev. Income Wealth* 55 (3), 661–685.
- Cowan, R., David, P.A., Foray, D., 2000. The explicit economics of knowledge codification and Tacitness. *Ind. Corp. Chang.* 9, 211–253.
- Dean, J., 2019. *The Deep Learning Revolution and its Implications for Computer Architecture and Chip Design*. Working Paper, Google Research. <https://arxiv.org/ftp/arxiv/papers/1911/1911.05289.pdf>.
- Domini, G., Grazi, M., Moschella, D., Treibich, T., 2022. For whom the bell tolls: The firm-level effects of automation on wage and gender inequality. *Res. Policy*. ISSN: 0048-7333 51 (7), 104533. <https://doi.org/10.1016/j.respol.2022.104533>.
- Duch-Brown, N., Gomez-Herrera, E., Mueller-Langer, F., Tolan, S., 2022. Market power and artificial intelligence work on online labour markets. *Res. Policy* 51 (3). <https://doi.org/10.1016/j.respol.2021.104446>. ISSN 0048-7333.
- European Commission, 2019. *The Impact of Digital Transformation on EU Labour Markets*. EC. April 2019. https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=58412.
- European Commission, 2022. *European Data Governance Act*. <https://digital-strategy.ec.europa.eu/en/policies/data-governance-act>.
- European Commission, 2023. *Data Act*. <https://digital-strategy.ec.europa.eu/en/policies/data-act>.
- Feigenbaum, J., Gross, D., 2022. Answering the call of automation: labor market adjusted to the mechanization of telephone operation. In: NBER Working Paper 28061. National Bureau of Economic Research, Cambridge, MA. <https://www.nber.org/papers/w28061>.
- Felten, E., Raj, M., Seamans, R., 2019. *The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization*. Working Paper, NYU Stern School of Business. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3368605.
- Fernandez Macias, E., Urzi Brancati, M.C., Wright, S., Pesole, A., 2023. *The Platformisation of Work*, EUR 31469 EN. Publications Office of the European Union, Luxembourg. <https://doi.org/10.2760/801282>, 2023, ISBN 978-92-68-01661-9. JRC133016.
- Foray, D., Steinmueller, W.E., 2003. The economics of knowledge reproduction by inscription. *Ind. Corp. Chang.* 12, 299–319.
- Foray, Dominique, 2004. *The Economics of Knowledge*. MIT Press.
- Fumagalli, E., Rezaei, S., Salomons, A., 2022. OK computer: Worker perceptions of algorithmic recruitment. *Res. Policy*. ISSN: 0048-7333 51 (2), 104420. <https://doi.org/10.1016/j.respol.2021.104420>.
- Furman, J., Coyle, D., Fletcher, A., Marsden, P., McAuley, D., 2019. *Unlocking Digital Competition: Report of the Digital Competition Expert Panel*. HM Treasury. www.gov.uk/government/publications.

- Genz, S., Gregory, T., Janser, M., Lehmer, F., Matthews, B., 2021. "How do Workers Adjust When Firms Adopt New Technologies?" Working Paper, ZEW – Centre for European Economic Research Discussion Paper no. 21-073. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3949800.
- Goldfarb, A., Tucker, C., 2019. Digital economics. *J. Econ. Lit.* 57 (1), 3–43.
- Goos, M., Manning, A., Salomons, A., 2014. Explaining job polarization: routine-biased technological change and offshoring. *Am. Econ. Rev.* 104 (8), 2509–2526.
- Graef, I., Prüfer, J., 2021. Governance of data sharing: a law & economics proposal. *Res. Policy* 50 (9). <https://doi.org/10.1016/j.respol.2021.104330>.
- Gravey, V., Savona, M., Wincott, D., Winters, A., 2023. Trade-offs in trade policy. What the public thinks and how they think about them. In: Centre for Inclusive Trade Policy. Briefing Paper 3. April 2023. <https://citp.ac.uk/publications/trade-offs-in-trade-policy-what-the-public-thinks-and-how-they-think-about-them>.
- Hoffmann, M., Nurski, L., 2021. What is holding back artificial intelligence adoption in Europe?. In: Policy Contribution 24/2021, Bruegel.
- Igna, I., Venturini, F., 2023. The determinants of AI innovation across European firms. *Res. Policy* 52 (2). <https://doi.org/10.1016/j.respol.2022.104661>.
- Li, D., Raymond, R.L., Bergman, P., 2020. Hiring as exploration. In: NBER Working Paper 27736. August 2020. <https://www.nber.org/papers/w27736>.
- van Lieshout, M., 2015. The value of personal data. In: Camenisch, Jan, Fischer-Hubner, Simone, Hansen, Marit (Eds.), *Privacy and Identity Management for the Future Internet in the Age of Globalisation*, 2015. Springer Verlag, London, pp. 26–38.
- Lynskey, O., 2015. *The Foundations of EU Data Protection Law*. Oxford University Press.
- Mahieu, R., 2021. The right of access to personal data: a genealogy. In: *Technology and Regulation*, 2021, pp. 62–75. <https://doi.org/10.26116/techreg.2021.005>.
- Mansell, R., Steinmueller, E.W., 2020. *Advanced Introduction to Platform Economics*. E. Elgar Advanced Introduction Series.
- Mateos Garcia, J., Klinger, J., Stathopoulos, K., 2024. A Narrowing of AI Research? *Research Policy* (forthcoming).
- Mokyr, J., 2005. Long-term economic growth and the history of technology. In: *Handbook of Economic Growth*, Vol. 1, pp. 1113–1180. Part B, 2005. <https://www.sciencedirect.com/science/article/abs/pii/S1574068405010178>.
- Organization for Economic Cooperation and Development (OECD), 1992. *Proposed Guidelines for Collecting and Interpreting Technological Innovation Data*. OECD and Eurostat, Paris.
- Organization for Economic Cooperation and Development (OECD), 2019. *OECD AI Principles Overview*. <https://oecd.ai/en/ai-principles>.
- Organization for Economic Cooperation and Development (OECD), 2019. *Measuring the Digital Transformation: A Roadmap for the Future*. OECD Publishing, Paris. <https://doi.org/10.1787/9789264311992-en>.
- Prytkova, E., Petit, F., Li, D., Chaturvedi, S., Ciarli, T., 2023. Revealing semantics: exposure of industries and occupations to emerging STI areas. In: *Pillars Working Paper* (forthcoming).
- Rolf, S., O'Reilly, J., Meryon, M., 2022. Towards privatized social and employment protections in the platform economy? Evidence from the UK courier sector. *Res. Policy* 51 (5). <https://doi.org/10.1016/j.respol.2022.104492>. ISSN 0048-7333.
- Rotolo, D., Hicks, D., Martin, B.R., 2015. What is an emerging technology? *Res. Policy* 44 (10), 18271843.
- Savona, M., 2019. The Value of Data: Towards a Framework to Redistribute It. *SPRU Working Paper Series (SWPS)*, 2019-21, pp. 1–22. ISSN 2057-6668. Available at: www.sussex.ac.uk/spru/swps2019-21.
- Savona, M., 2020. Governance models for redistribution of data value. *VOX CEPR* (17 Jan 2020).
- Savona, M., Ciarli, T., Steinmueller, W.E., Vannuccini, S., 2022. The design of digital automation technologies. Implications for the future of work. *CESifo EconPol Forum* 23 (5), 4–10.
- Specht, L., Zerbst, J., 2018. Considering the relationship between the civil law treatment of data and data protection law in Germany. *J. Intell. Prop. Law Practice* 13 (6), 504–512. <https://doi.org/10.1093/jiplp/jpy061>. June 2018.
- Tucker, C., 2019. Digital Data, Platforms and the Usual [Antitrust] Suspects: Network Effects, Switching Costs, Essential Facility. *Rev. Ind. Organ.* 54 (4), 683–694.
- Webb, M., 2020. The Impact of Artificial Intelligence on the Labor Market. Working Paper, Stanford University. https://www.michaelwebb.co/webb_ai.pdf.
- Weil, D., 2014. *The Fissured Workplace*. Harvard University Press.
- Westin, A.F., 1968. "Privacy And Freedom", 25 Wash. & Lee L. Rev. 166. Available at: <https://scholarlycommons.law.wlu.edu/wlulr/vol25/iss1/20>.
- Wood, A., 2021. *Algorithmic Management*. European Commission, JRC Technical Report. <https://joint-research-centre.ec.europa.eu/system/files/2021-05/jrc124874.pdf>.
- Zander, T., Steinbruck, A., Birnstill, P., 2019. Game-theoretical model on the GDPR: market for lemons? *J. Intell. Prop. Inf. Technol. E-Commerce Law* 10 (2), 200–208.

Maarten Goos^a, Maria Savona^{b,*}

^a *Utrecht University School of Economics, Netherlands*

^b *Science Policy Research Unit, University of Sussex, United Kingdom of Great Britain and Northern Ireland and Department of Economics and Finance, Luiss University, Italy*

* Corresponding author.

E-mail address: M.Savona@sussex.ac.uk (M. Savona).