



Technequality

Understanding the relation between technological innovations and social inequality

Deliverable D1.2

Scenarios for the Impact of Intelligent Automation on Work

Version 2.0

12-06-2020





Funded by the Horizon 2020
Framework Programme of the
European Union

Deliverable	:	1.2
Title	:	Scenarios for the Impact of Intelligent Automation on Work
Author(s)	:	M Levels ¹ , M Somers ² , M-C Fregin ³
Reviewer(s)	:	R Montizaan, P. Summerton
Document type	:	R
Document level	:	PU
Date	:	12 June 2020
Status	:	complete

TECHNEQUALITY partners

ROA Universiteit Maastricht

TiU Stichting Katholieke Universiteit Brabant

UOXF The Chancelor, Masters and Scholars of the University of Oxford

CE Cambridge Econometrics Ltd.

SOFI Stockholms University

WZB Wissenschaftszentrum Berlin für Sozialforschung GGmbH

EUI European University Institute

TU Tallinn University

¹ Research Centre for Education and the Labour Market, Maastricht University, Tongersestraat 49, 6211 LM, The Netherlands. E-mail: m.levels@maastrichtuniversity.nl

² Research Centre for Education and the Labour Market, Maastricht University, Tongersestraat 49, 6211 LM, The Netherlands. E-mail: meline.somers@maastrichtuniversity.nl

³ Research Centre for Education and the Labour Market, Maastricht University, Tongersestraat 49, 6211 LM, The Netherlands. E-mail: m.fregin@maastrichtuniversity.nl



Description of deliverable (100 words)

Deliverable 1.2 provides 8 scenarios for how current technological innovations will affect work, by examining three key variables: speed of innovation in key fields, speed of adaptation, and impact on work tasks.

*Time present and time past
Are both perhaps present in time future,
And time future contained in time past.
If all time is eternally present
All time is unredeemable.
What might have been is an abstraction
Remaining a perpetual possibility
Only in a world of speculation.
What might have been and what has been
Point to one end, which is always present.*

T.S. Elliot – Burnt Norton (1936)



Contents

1.	Introduction.....	6
1.1	Approach.....	8
2.	Potential of key technologies and innovation speed	11
2.1	Intelligent automation: potential of key technologies	11
2.1.1	Artificial intelligence	11
2.1.2	Robotics	14
2.2	Innovation speed: eliminating engineering bottlenecks	16
2.2.1	Engineering bottlenecks in artificial intelligence	17
2.2.2	Engineering bottlenecks in robotics	19
2.2.3	Potential outcomes.....	20
3	Adoption and diffusion	21
3.1	Understanding adoption.....	21
3.1.1	Characteristics of innovations	22
3.1.2	Characteristics of organisations	23
3.1.3	Context.....	24
3.2	Adoption rate: potential outcomes.....	25
3.2.1	Further assumptions on diffusion and adoption	25
3.2.2	Potential outcomes.....	28
4.	Channels through which technology affects work.....	29
4.1	Intelligent automation can substitute for human labour	29
4.2	Intelligent automation can augment human labour	31
4.2.1	Compensation through complementarities	31
4.2.2	Compensation through lower costs.....	32
4.2.3	Compensation mechanism via increase in incomes	33
4.2.4	Empirical evidence	33
4.2.5	Impact: potential outcomes	34
5	Scenarios for the future of work.....	36
5.1	Acute disruption	39
5.2	Incremental automation	40
5.3	Delayed disruption.....	42
5.4	Slow substitution	43
5.5	Abrupt volatility	44
5.6	Controlled adjustment.....	45
5.7	Delayed substitution	46
5.8	Gradual substitution	47
6	Discussion	48
	References.....	50



1. Introduction

Recent headlines regularly suggest that intelligent automation may usher in an age of technological unemployment. Widespread anxiety over a jobless future induced by automation and technological progress is not a recent phenomenon. All industrial revolutions so far were accompanied by recurring concerns that machines were going to eliminate a substantial number of jobs. Already during the first Industrial Revolution, a group of British textile artisans – better known as the Luddite Movement (1811-16) – revolted against the increased adoption of power looms and mechanical knitting frames by destroying the textile machinery out of fear of job losses. The view that technological progress can have lasting negative effects on working conditions and overall employment was underscored by famed economists including Karl Marx and David Ricardo in the 19th century. In the last century, John Maynard Keynes introduced the term “technological unemployment” which he defined as “unemployment due to our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour” (Keynes, 1933, p. 3). While historical evidence does generally not support fears of machines substituting for human labour in the past (Autor, 2015), recent advances in digital technologies and robotics have fed concerns about massive job losses (Hogarth, 2017; Pew Research Centre, 2017). The performance of many tasks that were considered beyond the potential of computer-assisted technologies are now within the scope of what computers can do (Brynjolfsson & McAfee, 2014; Ford, 2015; Frey & Osborne, 2017). This has led many to believe that “this time may be different.”

Predicting the impact of technological innovations in the field of intelligent automation on labour markets and inequalities in European societies is not straightforward. It is by no means given – or even plausible – that current technological and economic trends can be sensibly extrapolated into the future. While we may assess expectations of the potential impact of different technologies on the world of work, it is quite uncertain if, how, and when technologies will affect the demand for labour. This is partly inherent to uncertainty about the potential of technological innovations. Past technological and economic trends cannot be simply extrapolated into the future. The current wave of technological innovations, and in particular the combination of innovations in the fields of robotics, machine learning, and quantum computing, have the potential to be disruptive. For example, the potential of machine learning (ML) to learn complex patterns may be revolutionized by developments in quantum computing, which would tremendously reduce computing time and enable learning at an unprecedented speed. Moreover, ML and related technologies are considered “general purpose technologies” with the potential to affect the entire or at least large parts of the economy and compared to electricity.



Existing empirical studies do not provide a sufficient answer to how technological developments will affect the world of work with respect to *quantitative* outcomes (job destruction and job creation) but also *qualitative* changes in work and employment (de-skilling, up-skilling and the evolvement of new task profiles of existing jobs and “new collar work”). Several studies predicting the extent to which jobs are susceptible to automation are shaping the current debate about the labour market implications of technological progress. Estimates of the share of jobs that are at high risk of automation in the near future range from 9 to 47 percent (Artzn, Gregory & Zierahn, 2016; Frey & Osborne, 2017; Hawksworth, Berriman & Goel, 2018; Manyika et al., 2018; Nedelkoska & Quintini, 2018). The estimates produced by the aforementioned studies rest on an expert assessment of the type of tasks that are (still) difficult to automate given the current state of technology. Based on the task composition of jobs, such so-called engineering bottlenecks are used to assess the risk of automation for occupational groups. Despite the fact that the studies seeking to estimate the automation risk of jobs stress the caveats of their analyses, their estimates have repeatedly been interpreted as an indicator for actual job losses. However, the fact that some job tasks are automatable does not automatically translate into actual job losses or rising unemployment levels. This is because many factors (e.g., price and access to technology, legislation, availability of training data and also managerial practices and culture) constrain the adoption and diffusion of technologies (Habakkuk, 1962). Moreover, these studies assess to what extent it is technically feasible to *substitute* machines for human labour, but disregard the potential of machines to create jobs and to *complement* workers performing tasks that cannot be automated.

This report aims to contribute to our understanding of the potential implications of technological change and provides plausible future scenarios of how the recent wave of technological innovations will affect labour markets in EU countries. Uncertainty regarding technological potential and how the market for these technologies will react to the availability of new techniques, make it infeasible to predict with full certainty how the future will unfold. In such cases, scenario studies are useful tools. The goal of scenario studies is to analyse future scenario's by exploring and describing alternative potential outcomes, given certain crucial variables. To achieve this goal, we take the reader on a journey into the near future, describing eight possible scenarios from the perspective of the years 2025 and 2035, looking back on what could possibly have happened between today and 2025 or 2035 respectively to achieve these scenarios. This thought experiment allows us to pinpoint the different developments that can take place in the near future, exploring possible scenarios and potential outcomes along with respective policy implications. However, scenario studies are not prognoses or forecasting, and neither aim nor claim to predict the future. Rather, they serve as tools for understanding how various important variables may work to shape



that future. The eight different scenarios that we describe serve as thought-provoking tools that can help academics and policy makers to think about potential responses to intelligent automation.

The scenario's in this study also are related to other Technequality deliverables. For example, they form the theoretical underpinning for scenario's considered in D1.4, in which we estimate the actual impact of automation on jobs. In D1.4, we also quantify potential government responses to these scenario's, which is something that we do not do in the current study. The scenario studies also provide the basis for a number of Technequality transfer publications targeted at policy makers, notably those related to D7.1. As part of this deliverable, we published a Policy Report, a Policy Brief, and several Technequality Factsheets that are informed by the analyses of scenarios from this report.

The scenarios we developed are not intended as a normative interpretation of potential outcomes. The scenarios should be considered objective explorations of possible outcomes. It is important to note that the scenarios are not mutually exclusive. The future of work will most likely be characterized by a combination of various scenarios that unfold simultaneously, with some scenarios being more plausible in certain economic sectors and countries, and less in others.

The remainder of this whitepaper is structured as follows. The next section describes our approach towards analysing future scenarios. Sections 2-4 identify and analyse the factors that we deem most crucial in determining the impact of the current wave of innovations on the workforce. Section 2 focusses on technological factors and discusses the automation potential of key technologies. Moreover, Section 2 provides insight into the speed with which innovators are able to overcome bottlenecks in various technological domains. Section 3 and 4 concentrate on non-technological variables shaping the future of work. In Section 3, we describe the factors determining the actual adoption and diffusion of technological innovations. Section 4 discusses through what channels technological change can either lead to a displacement of workers or to an increase in the demand for labour. In Section 5, we describe eight plausible scenarios for the future world of work and discuss what implications each scenario entails for different groups of workers. In addition, we provide possible policy responses to each scenario. The final section concludes.

1.1 Approach

To develop potential outcome scenarios for the impact of intelligent automation, we identify key factors that we consider crucial for shaping the future of work. The importance of these factors is determined by an evaluation of past impact and importance, as well as considerations about future impact. To determine this, we built on research of scientific literature and grey literature, as well



as transcripts from the Technequality Expert Meeting on the Impact of Automation, organised April 24 and 25 in Amsterdam.⁴

This procedure led to the identification of three key explanatory variables. The first (technological) dimension concerns the potential and speed of technological developments. First, we start by explaining what the current generation of technological innovations that enable intelligent automation can and cannot do by drawing upon recent developments in engineering sciences. As the extent of automation is also determined by whether certain engineering bottlenecks will be overcome, we shed light on the pace at which these bottlenecks are likely to be solved. Second, we describe which factors affect the actual adoption of technologies. Third, because machines do not merely affect the workforce through labour substitution, we discuss through which (economic) channels technological innovations might decrease, but also increase labour demand. It is important to note that we consider the mechanisms through which technological change affects work as one of the explanatory variables. For example, whether innovations have a positive or negative net effect on employment depends on whether technologies substitute or complement human labour. In turn, the substitution or augmenting effect of technological change depends on other factors including labour costs, complementarities, and income effects. Hence, it is the combination of technological potential, the actual adoption of technologies and whether innovations predominantly complement or substitute human labour that finally determines what the future of work will look like. The outcome variable of interest is the actual impact of these variables on work, both in quantitative as well as qualitative terms.

We then describe eight potential future scenarios. These scenarios provide useful heuristics for understanding the implications of developments and potential frameworks for policy responses.

⁴ A description of this meeting is provided in the appendix.



Crucial variables

- 1. Potential of key technologies and innovation speed**
- 2. Adoption and diffusion rate**
- 3. Impact on work**



2. Potential of key technologies and innovation speed

The impact of intelligent automation on the world of work is determined by a number of technological factors. Technological factors concern the range of tasks that can be performed by machines and the speed with which innovators are able to overcome engineering bottlenecks in various technological domains. It is important to note that developments in various technological fields cannot be considered separately, since lack of progress in one field may be a crucial bottleneck in another. For example, engineering, training, and scaling algorithms are necessary conditions for unlocking the potential of artificial intelligence (AI). However, for machine learning algorithms to reach its full potential and to turn into AI, computing speed must increase, and memory possibilities need to be improved. Although quantum computing promises an astronomical increase in computing power and progress is being made, engineers still have to solve some of the most important technical problems (DiVicenzo, 1995; Arute et al., 2019).

Below, we first describe the technologies that are crucial for intelligent automation and we discuss their current potential (2.1). We then describe how we define innovation speed as a crucial factor for future scenarios, which depends on market forces and the ability to overcome crucial bottlenecks. We identify some main bottlenecks in the various technological fields (2.2).

2.1 Intelligent automation: potential of key technologies

2.1.1 Artificial intelligence

At the core of development of intelligent automation lies the capacity of machines to perform and learn tasks with a limited need for human interference. AI refers to systems that behave “intelligently” by analysing their environment and choose the best action to take – with some degree of autonomy – to achieve specific goals (AI-HLEG, 2019; Gesing, Peterson & Michelsen, 2018). AI-based systems can be entirely software based (e.g. spam filters, instant machine translation, search engines, voice and facial recognition systems) or embedded in hardware devices (e.g. warehouse robots, self-driving cars). AI systems have three main features: the ability to perceive, reason and decide, and act (AI-HLEG, 2019). For an AI system to perform tasks with some degree of autonomy, it must be able to perceive data present in the environment in which it operates through sensors (e.g. a robot apple picker sees and isolates individual apples from trunks and twigs through a camera). AI systems process the received data and decide on which action to take given a pre-specified goal (e.g. decide whether an apple is ripe or not). Next, the action is performed by possible manipulation of the environment (e.g. picking an apple or not). AI systems



are particularly capable to learn how to solve problems that cannot be precisely specified in a well-defined set of rules. The performance of such tasks is enabled through a set of techniques that follow the machine learning (ML) approach.

The simple idea of ML is to programme computers to learn from example data in cases where formal procedural rules are unknown. Consider, for example, what it takes to let a computer identify a lavender plant in a flower auction warehouse. We could programme a computer to select this type of plant based on its colour and shape. However, there are likely to be other plants that are similar in terms of their colour and shape, such as the wild sage plant. While we could attempt to describe lavender plants in finer detail, the amount of complexity increases tremendously. Instead, machine learning enables us to program a computer to become proficient at accomplishing a task autonomously by exposing it to examples of tasks that were performed successfully by others. Through a process of exposure, training, and reinforcement, machine learning algorithms can infer how to master tasks that are otherwise difficult to program into a well-defined set of rules.

While recent advances in the field of machine learning expand the potential for intelligent automation, they are not equally suitable for all tasks. To understand what tasks are automatable through machine learning, it is helpful to understand what types of learning problems comprise the vast majority of applications of machine learning. Depending on the degree of human participation and the real-time integration of the feedback on the performance of the system, various training methods can be distinguished. The basis, however, is always a learning problem. The most wide-spread approaches of machine learning are *supervised learning*, *unsupervised learning*, *semi-supervised learning*, and *reinforcement learning* (AI-HLEG, 2019). Below, we discuss what the potential of each of these machine learning techniques is in terms of task automation. Moreover, we explain which conditions must be met - or equivalently, which engineering bottlenecks must be solved - to successfully apply these techniques to real world problems.

Most successful commercial applications of machine learning draw on a technique called *supervised learning*. Supervised learning is particularly useful in cases where an input A needs to be transformed into an output B (Ng, 2016). For example, a computer can learn by studying examples of correct input (e.g. medical record) and output (e.g. the likelihood that a patient has cancer) combinations. In supervised learning software, a machine learns a task under the supervision of a human, typically on the basis of fully labelled datasets. Fully labelled implies that each example in the training dataset

is tagged with the answer or outcome the algorithm should come up with on its own (Salian, 2018). For instance, we might want to train an algorithm to classify whether an email message (input) is spam or not (output) by providing it with a large collection of emails with their correct labels (i.e. spam or no spam). To successfully apply supervised learning in practical settings, the task to be learned must be fully specifiable, along with a performance metric, and training experience (Salian, 2018). Relying on large databases of so-called ground truth (optimal decision) – specifically, a large set of examples of labelled objects – a machine learning algorithm can attempt to statistically infer what the relationship is between some input and output combination.

In an *unsupervised learning* model, unlabelled data is provided to the algorithm that attempts to extract features and patterns from the data autonomously (Salian, 2018). In unsupervised learning problems, datasets are being explored without a concrete pre-specified target function that needs to be learned. Unsupervised learning can be applied to cluster website visitor records into groups of visitors with similar characteristics or behaviour, to help understand what types of people visit a specific website. Because this technique is often prone to biases as there is no “ground truth” in the data that could serve as baseline, this type of machine learning has led to fewer commercial applications. *Semi-supervised learning* takes a middle ground. It uses a subset of labelled data in combination with a larger set of unlabelled data, improving the accuracy of unsupervised learning (reference). A common practical example is the analysis of medical images and CT scans. The technique is also useful for training robots to manage inventory in a warehouse.

The most advanced learning technique at the moment is *reinforcement learning*, where an algorithm is trained with a reward system, providing feedback when it performs the best action in a particular situation (Salian, 2018). As in supervised learning, the learning algorithm is required to learn a specific target function. In contrast to supervised learning, the self-teaching system learns through trial and error and aims at achieving the best outcome. For example, reinforcement learning is used when a robot learns how to walk. Imagine a robot trying a large step forward that falls. The outcome of a fall, combined with a big step, becomes a data point that the reinforcement learning system learns from. As the feedback was negative (the outcome was undesired), the system adjusts its action by trying to take a smaller step. The learning algorithm’s goal is to choose actions that leads to the desired outcome or reward. In order to successfully learn how to accomplish a task, the learning algorithm must train itself by making a great amount of iterations in its actions. This means it can in principal independently master a wide variety of tasks. As a consequence, Andrew

Ng (2016) predicts that “If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future. (Ng, 2016).

2.1.2 Robotics

A robot is a mechatronic device, that has sensors and actuators, and that can freely move around a physical space to a certain degree. Crucially, robots can autonomously perform a variety of tasks, without the tasks being programmed in advance. Typical modern robots have a torso or frame, mechanical and electronic means to move (e.g. omni wheels, motors and power supplies) the ability to manipulate the environment (e.g. mechanical arms and hands), a perception system (potentially including lidar technology, laser range finders, camera, and data processing capacity), computer systems, and a human-machine interface (Scholtes, 2019). Robotics can also be considered as AI that acts in the physical world, also defined as *embodied AI* (AI-HLEG, 2019).

The more autonomous a robot, the better capable it is to perform non-routine tasks. Robots are designed to operate in three types of physical environments, each demanding a different level of autonomy. The most limited environment assumes a *closed world*. Robots in closed world environments perform known tasks in known environments. The world in which they operate is characterised by extremely low variation, which is controlled and fenced off, and everything is engineered. These robots are not adaptable and they cannot reconfigure. If one robot in a production line breaks down, the whole line stops operating. As such, they require an extensive safety system. Closed world robots are limited in capacity, have little to no perception systems, and are 100% pre-programmed and engineered. They do however have a high reliability. To be economically feasible, closed world robots must perform high volumes of tasks. Such robots have been around for decades and will remain operational in the future. A good example of robots working in a closed world scenario are the robots that are currently operational in logistic supply centres, such as the drive units, palletizers and robo-stows that operate with humans at Amazon fulfilment centres (Scholtes, 2019; Amazon, 2019).

Robots that are designed to operate under a *semi-open world* assumption are already more versatile. They can perform known tasks in known environments, with a known nominal state of the world, but with a higher variation. The level of variations in the semi-open world are known, as are the potential tasks and environment conditions. These robots are able to adapt and account to known variation. The advantage over the closed world robots is that it is not known where, when, or in what order events take place, such that robots operating in a semi-open world can deal with a high



diversity of unique events. Such robots have perception systems that can account for known variation, and do not rely so much on programming. They are able to (re)configure at the level of skills. To be economically feasible, such robots must perform a medium to high volume of a variety of tasks. These types of robots are currently being produced for various industries, and will make up the better part of the market for the next 10 years. Current examples include fruit picking machines, cow milking robots, and many types of cobots (short for “collaborative robots”; Scholtes, 2019; Peters, 2019).

The most versatile robots are designed to function under an *open world* assumption. These robots can perform known tasks, in unknown environments, encounter and handle unknown objects, and engage in unknown interactions. Their software allows them to deal with a high variation of products and operating environments and conditions. They are highly adaptable and able to reconfigure their own actions depending on the task or operating environment. Such robots require semantic and context awareness, the ability to account meaning to events, and the ability to explain behaviour and understand causality. They must have over-dimensional perception systems that allow them to interpret and adapt to unknown variation. Robots like these cannot be pre-programmed but must rely on self-learning and self-configuration at the level of tasks. Robots like these can be used for unique tasks. Such robots do not yet exist but are likely to be developed within the next 10-30 years (Scholtes, 2019). Several bottlenecks must be overcome before robots can operate in open world environments. A good example are autonomous driving vehicles. While they must have the potential to operate fully in the uncertain environment of traffic, current autonomous cars can operate somewhere in between the semi-open and open world environment. The systems are not yet robust to respond adequately to high level of variation (such as weather, traffic conditions, dumped cars) and still require detailed road maps (Scholtes, 2019).

While robots are increasingly capable of performing new tasks, many tasks that seem simple to humans are far beyond the reach of machines. This has been called Moravec's (1988) paradox: "it is comparatively easy to make computers exhibit adult level performance on intelligence tests or in playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility." However, as the development of AI progresses, a new generation of robots is emerging and partly even ready-to-use in industrial production. Using AI-based high-performance sensor technology, these robots can perform fine motor tasks and operate in open spaces at an increasing rate. Moreover, robots increasingly work with humans in physical collaboration as cobots. These collaborative robots are equipped with self-optimising algorithms

that allow them to learn from their human colleagues. Hereby, the machine is supposed to reduce the likelihood of repeated mistakes. For example, if the desired action to a proposed problem is uncertain as no clear evidence is available, the machine could request confirmation or indication of a label by the expert. A prerequisite is that the actions and responses of the machine should be comprehensible to humans.

2.2 Innovation speed: eliminating engineering bottlenecks

From Section 2.1, we derive a number of plausible future developments concerning the potential of technologies. These are reasonable assumptions – not deterministic predictions – about the future that we deem plausible given current developments.

1. Algorithms will eventually enable machines to perform both routine and non-routine tasks at a higher proficiency level than humans, in multiple economic sectors.
2. Industrial robots will increasingly be able to operate in a larger number of semi-closed world environments, and eventually be able to operate in an open world environment.
3. Machines will increasingly be able to learn autonomously based on reinforcement learning

The speed of innovative developments is largely dependent on the elimination of major engineering bottlenecks. This is well illustrated by past trends in the field of AI. While relevant research has been ongoing since the 1950s and the term AI was already coined at the Dartmouth College conference in 1956, extensive periods of time witnessed relative slow progress. Thereafter, limited computing power and limited storing capacity ushered in the first AI winter, in which developments stalled and investments slowed down. In 1964, the first chatbot (Eliza) was developed at MIT, which renewed attention in AI research, culminating in the development of the first expert system by Feigenbaum between 1975 and 1982. After that, investment again dropped due to failure to solve complex real-life problems and achieve scale, and disappointing results. The 1997 performance of IBMs DeepBlue in beating Gary Kasparov at chess was a turning point, as was the 2011 Jeopardy! victory by IBM Watson and the development of Apple's Siri. Since then, AI developed with high speed, especially since 2015. This was enabled by largely increasing storage possibilities for Big Data and the increasing availability of powerful GPUs that not only made parallel processing ever faster but also cheaper (European Commission, 2018; Copeland, 2016). Parallel developments of growth in computing power, progress in self-learning algorithms, and Big Data have turned AI into one of the most important strategic technologies of the 21st century (European Commission, 2018, 2). Crucial breakthroughs are occurring at ever-faster pace. In 2015,



AlphaGo defeated human experts in the game of Go (Silver et al., 2016), while DeepMind learned to play 42 Arcade games through transfer learning (Mnih et al. 2015). In 2017, Libratus and Deepstack bested expert human poker experts (Brown & Sandholm, 2018; Moravcik et al. 2017), and AlphaGoZero learned to play Go without human intervention and managed to defeat human experts (Silver et al., 2017).

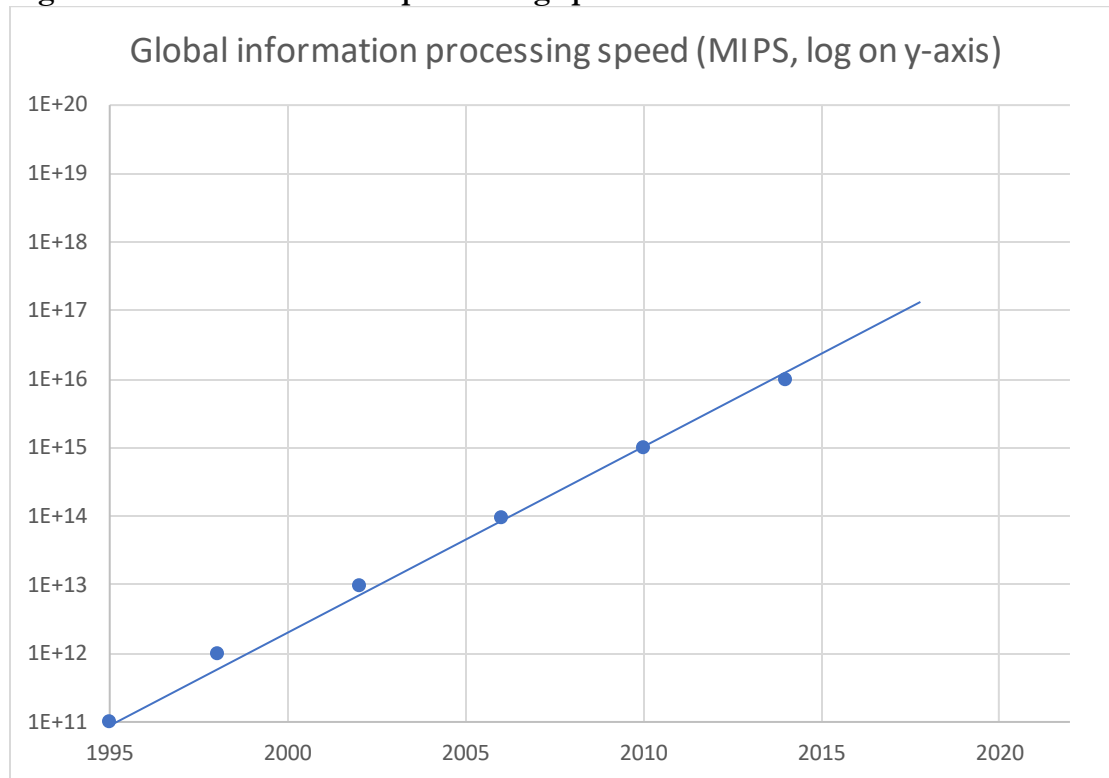
2.2.1 Engineering bottlenecks in artificial intelligence

While “general AI” – that is intelligent machines that have all human senses and reason – are still science fiction, “narrow AI” technologies are largely restricted to perform specific tasks as well as, or better than, humans can (e.g. read and analyse large volumes of texts) (Copeland, 2016). However, many bottlenecks remain. As described in Section 2.1.1., an important requirement for automating a task through ML is the availability of a large training dataset that provides examples of correct input (e.g. medical record) and output (e.g. diagnosis) combinations. Some ML techniques require thousands of data points to achieve acceptable performance in classification tasks and, in some cases, require millions of data records to perform tasks at a level that exceeds humans. It is estimated that a supervised deep learning algorithm will typically need a training data set containing around 5,000 labelled examples to perform a task at an acceptable level, while at least 10 million labelled examples are needed to match or exceed human level performance (Goodfellow, Bengio & Courville, 2016; Ng, 2016; Brynjolfsson, 2014; Brynjolfsson & Mitchell, 2017; Brynjolfsson, Mitchell & Rock, 2018). Training ML models are particularly demanding when complex and unstructured data types such as images, videos, audio or speech need to be processed (MGI, 2018). The low availability of ground-truth training data for many domains where this is now difficult to obtain, such as psychiatric diagnosis, hiring decisions, and legal cases is an important engineering bottleneck to broaden the automation potential of learning algorithms.

Additional key bottlenecks for further unlocking the potential of intelligent machines derive from limitations pertaining to hardware. As a data-driven technology, ML currently reaches its limits when large amounts of data have to be generated and learned from at the same time. Training ML models is the most computationally intensive task. At the same time, storage systems must be able to feed the ML model with a sufficient amount of data to optimize training performance. Here, developments are fast. Figures 1 and 2 illustrate that computing speed and data storage capacity have increased logarithmically for years (Waldrop, 2016). Moreover, in 2019, the Sycamore computer reportedly reached quantum supremacy, and performed computations at a fraction of the processor time used by the world’s fastest supercomputer (Arute et al., 2019).

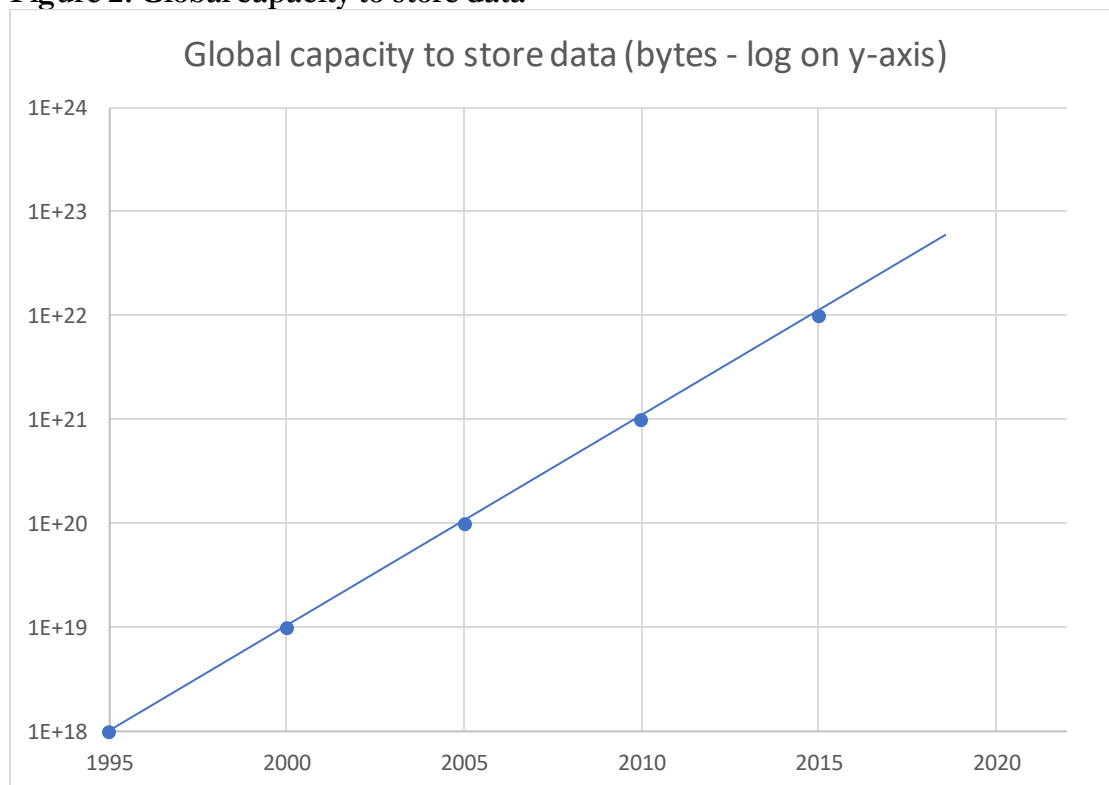


Figure 1. Global information processing speed



Source: Waldrop, M. M. (2016). *The chips are down for Moore’s law*. *Nature*, 530, 144–147

Figure 2. Global capacity to store data



Source: Hilbert, M., & López, P. (2011). *The World’s Technological Capacity to Store, Communicate, and Compute Information*. *Science*, 332(6025), 60–65.

2.2.2 Engineering bottlenecks in robotics

In the domain of physical tasks, machines are still not able to perform tasks requiring fine motor skills and to function in unstructured environments at human level performance (Yang et al., 2018). For example, sorting items for a customer's order from containers and placing them together in boxes is a challenging task from a machine-perspective. Items vary in size and stiffness and are initially stored together in a cluttered space (such as a warehouse shelf or in a fridge), making it challenging to identify and mechanically grasp objects.

To identify the desired objects, robots require sensors to assess colour and depth information. Improvements in sensors – along with falling prices – increasingly allow machines to operate in unpredictable and unstructured environments. This is mostly driven by developments in the telecom, gaming and automotive industries, that increase the availability of affordable IMU camera's (e.g. in smart phones), RGBD camera's (e.g. in the XBOX Kinect), ultrasonic sensors (e.g. parking assist sensors), and units for compact, low-energy, high performance computation (e.g. in smart phones). Another important driver of the price decline in sensors is the industry push for autonomous driving, that has spawned improved sensors such as LiDAR and Radar and stereo-cameras.

With respect to hardware challenges, the articulation ability of robot grippers need to be improved to match the flexibility and dexterity of human hands (Yang et al., 2018). The degrees of freedom defines the motion capabilities of mechanical devices. While human hands count up to 27 degrees of freedom, the degrees of freedom of most robot grippers is limited to 6 degrees of freedom. Consequently, robot grippers cannot articulate sufficiently to grasp most types of objects. Moreover, to manipulate objects, human hands are able to create the necessary friction interfaces between the object.

Developing and manufacturing hardware takes time and is expensive. The technologies are not as easily scalable as software deployment. It takes on average 3-5 years to develop an idea into a minimum viable product, even if funds are available. Production in this phase costs about EUR 5 million. Commercial applications require high level of reliability, so market integration will generally be done only after extensive testing. Therefore, we may expect developments in robotics innovations to evolve linearly rather than exponentially.



Engineering bottlenecks in the field of AI (see Section 2.2.1) are also critical for unlocking the automation-potential next-generation robots. The increased availability of big ground truth data sets, open source data sets, tooling, more efficient processing by GPU's and scalable parallel computation already led to improved image recognition (Scholtes, 2019). Further improvements in object recognizing and decision-making capacity on what action to take will depend on access to data as well as inexpensive and powerful computing hardware.

2.2.3 Potential outcomes

For reasons of simplicity, we limit ourselves to two potential counterfactual outcomes. Both assume that innovation in key technologies will continue, and that various public and private labs and R&D departments keep on working to overcome the remaining technical bottlenecks of key technologies for intelligent automation. They also assume that many of these bottlenecks will be overcome in the near future, and when bottlenecks are overcome in one field, there may be spill-over effects in another. Breakthroughs in visual pattern recognition software for example may help robots to navigate more effectively through semi-open worlds, and eventually through open-space scenarios.

The key variable for distinguishing scenarios is time. The pace with which innovation takes place, is generally unpredictable. Whereas some crucial trends plausibly evolve exponentially (e.g. the increase of processing speed), many of the crucial technologies will plausibly develop linearly and relatively slowly (e.g. the increasing availability of training data in many sectors). This suggests a **gradual innovation speed** as a plausible outcome. However, when engineering and technical bottlenecks are overcome, innovation may dramatically speed up and technologies may reach their potential faster. We therefore also consider an **innovation boom** as an interesting potential outcome.

3 Adoption and diffusion

The mere availability of technology does not automatically imply that innovations are actually adopted by industries and diffused across markets. A good example may be the implementation of machine learning algorithms and AI. These technologies already affected businesses in many economic sectors, but not quite as fast as many assume. Indeed, a large 2019 survey indicates that the vast majority of US companies are currently piloting AI implementation only on a small scale, ad hoc and in single business processes (Fountaine, McCarthy & Saleh, 2019). However, studies also indicate that the vast majority of EU businesses expect AI to significantly affect their value chain within the next three years, and are currently labouring to adopt HRM policies that will allow their companies to be ready for AI (Mercer, 2019).

3.1 Understanding adoption

So, what factors underlie the speed with which new technologies are diffused and adopted? Various models of technology diffusion all point toward economic, practical and sociological considerations that drive decision-making about the adoption of new technologies into business models and value chains. The classic and most well-known model for describing how innovations are adopted and diffused in networks is developed by Beal, Rogers, and Bohlen (1957) and refined by Rogers (1962). This model describes an adoption-curve of technologies as a Gaussian distribution. According to this model, 2.5% of the market are *innovators* who are the first to adopt an innovation. Somewhat later adoption is done by *early adopters* who make up 13.5% of the market, and the following 34% are characterised as the *early majority*. After this phase, 50% of the companies on the market have adopted the innovation. Thereafter, the *late majority* (34%) and *laggards* (16%) take to adopting the technology.

The model is essentially a heuristic that can help to understand how and at what rate technological innovations are adopted and diffused across markets. It was primarily intended to understand the diffusion of farming practices (Bohlen and Beal, 1957) and has since been applied to analyse the adoption and diffusion of diverse innovations, including ICT (*c.f.* Venkatesch et al., 2003), the Internet of Things (Ancarani et al., 2019), and mobile banking (Shaikh & Karjaluoto, 2015), but also for innovation of government policies (Berry & Berry, 2018). From this model we derive three main drivers of adoption and diffusion of innovations by firms, that we present here in a way to explain technology adoption by corporate actors with the express intent to maximise productivity and become or remain competitive on a market:

- 1) Characteristics of the innovation (here: technological innovations enabling intelligent automation)
- 2) Characteristics of actors (in this case: corporate actors)
- 3) Context (e.g. related to government regulation, market forces, or social pressure)

3.1.1 Characteristics of innovations

As discussed before, technologies like AI and robotics are increasingly able to perform routine and non-routine tasks historically performed by humans. But these capabilities alone are not sufficient for explaining why firms would invest in adopting them. When facing the decision to invest in new technologies, companies assess several aspects of these innovations. Greenhalgh et al. (2004) provide a systematic review of key characteristics that generally apply to the adoption of innovations. In general, innovations that offer a relative advantage in productivity or cost-effectiveness are more likely to be adopted. However, relative advantage is a necessary but not a sufficient condition for adoption. Innovations that are easily implemented, can be pilot-tested, and that have easily observable results are more likely adopted. Adaptability to current practices also increases the likelihood of adoption. Furthermore, risk, relevance for task performance, and knowledge requirements are also important considerations.

Similar considerations are made by whether firms expect to gain from investing in smart technologies. When it comes to AI, questions that guide decision-making around investment revolve largely around effects on competitiveness (see for example: Gerbert, Justus & Hecker, 2017; Mohanty & Vyas, 2018). When deciding how to invest, firms aim at exploring how investments in a certain technological innovation may shape firm's competitive advantage, relative to the current situation. For example, does the implementation of AI create more value and increase competitiveness? Can it satisfy customer needs more effectively? Will current processes be enhanced to make work flows more efficient, and save costs? Can current processes be altered to integrate innovations? Are the necessary preconditions for implementations met (e.g. regarding the availability of data flows to train the algorithms and the availability of human resources)?

And, of course, price is an issue. Firms are prone to evaluate if the potential gains justify the costs associated with investment in innovative technologies. For AI, most of the source material for creating algorithms is freely available through open sources, so here, companies interested in implementing AI into their value chain have to make a build or buy decision (*cf.* Gerbert et al., 2018). However, costs associated with implementation and training of staff may still be

considerable. For the adoption decision regarding industrial robots or fast computers, the purchasing price is an even more important consideration.

3.1.2 Characteristics of organisations

Some of the main bottlenecks for implementing new innovations are associated with characteristics of firms and organisations. The introduction of new technologies not only offers opportunities, but also brings along complex challenges for organisations. Organisational design theories emphasize that organisational structures can either facilitate or hamper the adoption of new products and processes (Mintzberg, 1979; Williamson, 1975). Examples of organisational properties that could slow down innovative behaviour include hierarchical decision-making structures and a lack of vertical integration in the production chain (Burns & Stalker, 1994; Teece, 1996). Although peer-reviewed studies do not exist, some white papers from the grey literature suggest the implementation of intelligent automation is also affected by organisational structures. For example, one key practical barrier to implementation of AI in firms involves cultural and organisational bottlenecks. These include problems of implementing AI in existing work-patterns, the potential to produce and analyse relevant data, organisational cultures, and the availability of human resources. Indeed, the problems to integrate AI into existing operating procedures prevents many early adopters from achieving results at scale (McKinsey, 2018; Gartner, 2019).

Another challenge remains labelling data. A critical step is to fit the AI approach to the problem and the availability of data. Since these systems are “trained” rather than programmed, the various processes often require huge amounts of labelled data to perform complex tasks accurately. To enable e.g. autonomous vehicles to drive in an open world, all critical objects must be labelled under all conditions of weather and light. Obtaining such large data sets can be difficult. In some domains, they may simply not be available, but even when available, the labelling efforts can require enormous human resources (Chui, Maniyaka & Miremadi, 2018). The organisation must be able to produce such labelled data; not all organisations are structured for this data production, and often core processes are not easily captured in data.

Cultural barriers are also important in determining the success of the implementation of AI (Fountain et al., 2019). Successful innovation requires broad acceptance and embeddedness in organisations. For example, the recommendations generated by AI systems can assist workers to arrive at better answers than either workers or machines could obtain independently. However, this will only come about if workers trust the answers or suggestions provided by machines. Moreover, this approach requires that the decision-making process should be supported by the AI system and the traditional top-down approach needs to be decentralized. The involvement of end



users in the design of applications will increase the acceptance of new technologies, and, thereby, make successful implementation more likely. Hence, the extent to which workers are used to use a software to guide their work (doctors are not at all used to be *guided* by software while talking to patients – by contrast, call centre agents are) as well as the overall readiness of the organisation for the use of AI can affect both quantitative and qualitative outcomes.

Another main issue is the availability of necessary human resources. Working with machines requires a specific skill set that may differ from a skill set traditionally associated with a certain job. Building and maintaining the technology required for intelligent automation may require technical staff that is proficient in new technologies. Indeed, the absence of qualified staff is one of the main bottlenecks for implementation of AI by organisations. AI needs to be implemented, trained, customized to the business context and the firm's data, and then be promoted and used. The required skills are currently rare (Ng, 2016).

3.1.3 Context

The context in which firms operate is also extremely relevant for the pace of technology adaptation. Particularly, government regulation and relevant legislation, as well as market expectations, and pressure by stakeholders and interest groups.

Government regulation may stimulate or impede the adoption of technologies (Goolsbee, 2017). Robot taxes, for instance, are often thought to hamper investment (World Bank, 2019). Also, policies related to privacy might affect the rate of technological diffusion. Technology providers as well as companies that deploy smart technologies have to assure that ML applications are compatible with legal requirements. The potential of AI is critically dependent on data quality, which in turn strongly hinges on compliance with privacy regulations. Consequently, ethical concerns about privacy restricted data use might stifle investment in AI. Moreover, liability and accountability for algorithmic decisions taken have to be both technically feasible and compatible with the systems in place (Fraunhofer 2018, 6). Market expectations may also be crucial. In competitive markets, bandwagon effects play a role, specifically if a new technology is hyped and firms assume that competitors are investing in technology. The digital transformation not only involves technical and economic changes as the role of human labour is also subject to change. Unions are expected to ensure that sufficient (re-)training options are available to mitigate the negative effects of technological change on the workforce. To avoid social conflict, unions are likely to demand that gains from increased productivity and equitably distributed are reinvested in employees.

3.2 Adoption rate: potential outcomes

From these considerations, we may draw a number of conclusions about plausible future developments. These predictions are assumptions we deem plausible, but not beyond discussion. They do, however, form the foundation under our future scenarios.

1. The pace of diffusion of key technologies may differ between geographical regions, as different geographical regions represent different contexts.
2. The pace of diffusion of key technologies may largely differ between economic sectors, and variation within sectors between companies may be large.

3.2.1 Further assumptions on diffusion and adoption

The adoption rate of key technologies is the relative speed at which firms adopt key technologies. Following Rogers (1962) and common practice, we define the adoption rate of key technologies as the time it takes for a given percentage of the firms to adopt innovations in key technology. The expected adoption rate is determined by assumptions about the shape of the adoption curve. Various models are possible.

One often used model for predicting technological diffusion is the Gartner hype cycle. This is a graphic representation of technology adoption designed by Gartner. The methodology is intended to provide insight into how a technology or application will evolve over time (Gartner, 2019). The model distinguishes five stages:


- 1) Innovation Trigger: proof-of-concept and significant public attention to potential. Usable products do not yet exist, and commercial viability is unproven.
- 2) Peak of Inflated Expectations: Success stories gain further attention, but there are also many failures.
- 3) Trough of Disillusionment: Implementations of early products fail to deliver on the promise of the early success stories and public interest in the innovation wanes. In this stage, some technology products cease to exist. Further investment only if the early products satisfy early adopters.
- 4) Slope of Enlightenment: here, the added value of the innovation for firms become better understood and increasingly clear to more parties. The innovation becomes more widespread.
- 5) Plateau of Productivity: Mainstream adoption takes off.

The hype cycle is a much-used heuristic. It has been used for example to understand diffusion and adoption of AI (Columbus, 2019; Hildesheim, 2019). However, there is no clear empirical evidence to suggest that the stages in the model accurately describe the adoption of the technologies key to intelligent automation. Even if the adoption-curve would be shaped as the hype cycle suggest, the placement of the various key technologies on the curve is to a large extent arbitrary.

For reasons of simplicity, we therefore adopt the classic assumption about the shape of adoption curves (Rogers, 1962) and assume that the adoption of innovations follows an S-shaped curve, not unlike a logistic function. After initial slow adaptation, after a certain point in time, the diffusion of technology increases as the early majority jumps on the bandwagon. As the market share of the product reaches saturation, the pace of adoption slows down again, and then falls flat. Figure 3 suggests that the market adoption of industrial robots is now beginning to increase. If the assumption of a logistic curve is correct, the market penetration of these robots will increasingly speed up within the next 5-10 years.


Figure 4 is a projection by market research firm on emerging technologies Tractica (2019), that suggests that a similar trend may be expected for the market penetration of AI. According to this projection, the global artificial intelligence software market will grow significantly and exponentially in the next 10 years. Of course, it should be noted that these projections cover a broad group of technologies that can all be called AI, including machine learning techniques, natural language processing, and other technologies. This projection is predicated on the notion that organisational bottlenecks will eventually be overcome, and that this will allow the early adopters to achieve results at scale. Several analyses from the grey literature support the notion that hurdles remain. For example, McKinsey (2018) observes that AI has a huge potential in many sectors and that many companies have started to invest, but that few companies are ready to implement AI to achieve results at scale. Gartner also observes that AI adoption is only at the beginning of the curve (Miller, 2019).

Figure 3. Global sales of industrial robots (per 1000 units)



Source: International Federation of Robotics (2017). *World Robotics 2016 Industrial Robots*. Frankfurt am Main, Germany: IFR.

Figure 4. Artificial intelligence revenues (projection)



Source: Tractica (2019). *Artificial Intelligence Market Forecasts*. Boulder, CO: Tractica.



3.2.2 Potential outcomes

For the purposes of this study, and for reasons of simplicity, we limit ourselves to two potential outcomes. Both assume that market forces will compel companies to invest in new technologies, and that the diffusion of innovations will generally follow the logistic curve described in section 3.3.1. However, the crucial distinction is the time it takes for adoption to pick up speed.

The first is a trend of **slow adoption**. In this scenario, diffusion of innovations follows the exponential curve described in the previous section, but the organisational and cultural bottlenecks take relatively long and it takes at least a decade before early adaptors have successfully implemented intelligent automation and for the early majority begins to feel the market pressure to invest.

In contrast, the counterfactual **fast adoption** scenario assumes that organisational bottlenecks are overcome relatively fast, and that technologies are implemented at scale with early adaptors before 2029. In this scenario, bandwagon effects can further speed up diffusion of technologies in relevant markets.

4. Channels through which technology affects work

The third key variable that drives future scenarios pertains to how technological innovations affect actual work. From the literature, the effects of technology on the workforce appear ambiguous. New technologies are typically introduced to gain efficiency and reduce labour costs. However, a number of economic forces can compensate for the initial labour-saving effect of technological change. This section discusses the potential effects of technological change on (overall) employment and the mechanisms behind this relationship. Moreover, this section provides empirical evidence for the possible adverse effects of technological change on employment and its potential to counterbalance the initial displacement of jobs. This section starts with a discussion on the channels through which technological change can have labour-saving effects, which is followed by an overview of the countervailing forces through which technologies can have a labour-augmenting effect. Moreover, this section discusses the empirical evidence related to those mechanisms.

4.1 Intelligent automation can substitute for human labour ...

The production of goods and services in most industries involve the simultaneous accomplishment of a series of tasks. Skills (embedded in human labour) and technologies (embedded in machines or capital) can be considered competing inputs for the performance of various tasks (Autor et al., 2003; Acemoglu & Autor, 2011; Acemoglu & Restrepo 2018; Benzell et al., 2015; Susskind, 2017). According to recent labour market models in economics, firms choose the optimal allocation of workers and machines to tasks depending on the prices (equivalently, wages) and the relative productivity (i.e. comparative advantage) of these inputs in specific tasks. If machines become sufficiently cheap or sufficiently productive, then automation will lead to the direct substitution of machines for human labour in the performance of these tasks. This capital-labour substitution results in a displacement of workers from the tasks that are being automated.

Empirical work illustrates that the introduction of new technologies can adversely affect individual workers, at least in firms or industries in which the automation event takes place. The micro-level effects of automation are perhaps best understood by studying the impact on workers in firms where the automation originates: the automating firms. Using firm-level automation expenditure data across private sectors in the Netherlands over the period 2000-2016, Bessen et al. (2019) assess the impact of firm-level automation on a number of outcomes for incumbent workers as well as recent hires. The authors report that both incumbent workers and recent hires are more likely to leave their employer when the firms invest in automation technologies. The increased



probability of firm separation is already observable as from the first year in which automation is introduced and persists at least five years after the automation investment took place. While both incumbents and recent hires experience higher firm separation rates caused by automation, only incumbents suffer from a sizeable increase in the number of days annually spent in non-employment in the five years following the automation event. In contrast, recent hires enter new jobs rather smoothly. Automation events at the firm also translate into a substantial income loss for workers with longer firm tenure, while no such wage penalty is observed for recent hires. The income loss experienced by incumbents is almost entirely driven by non-employment spells followed by firm separation. These heterogeneous adjustment effects following job displacement might be explained by the fact that incumbent workers typically have accumulated more (firm-specific) human capital which depreciates in the light of automation. New hires, on the contrary, are more likely to be assigned to new job tasks when entering a new job (Raposo et al., 2015; Lefrance, 2003).

Similarly, Harrison et al. (2014) study the impact of process innovation on employment in a sample of firms in France, Germany, Spain and the United Kingdom. Process innovation implies that the same amount of output is produced with less labour inputs (Vivarelli, 2014). The authors find that process innovation causes a reduction in employment.

The direct displacement of workers from the production process can also be observed within automating industries. Using data for 32 industries in 19 OECD countries between 1970 and 2007, Autor and Salomons (2018) demonstrate that automation displaces workers and reduces labour's share of value added in industries where technological innovations are introduced.⁵ The micro-economic evidence on the effect of automation discussed here informs us that automation can *directly* cause a decline in the demand for labour in the firm or industry where it originates. It is important to stress that these findings do not imply that the adoption of automation technologies leads to labour displacement at the aggregate level. Direct displacements effects that are observed in the automating firm or industry can be offset by increases in employment and labour share elsewhere in the economy. Even within affected industries, automation technologies can spur labour demand (Acemoglu & Restrepo, 2018; Bessen, 2018; Brynjolfsson & Mitchell, 2017). In the next section, we turn to the countervailing effects through which technological progress can indirectly increase labour demand and lead to the reallocation of workers.

⁵ Autor and Salomons (2018) measure automation as industry-level changes in total factor productivity.



4.2 Intelligent automation can augment human labour ...

Macro-economic theory and empirical evidence illustrates that there is no direct link between automation and labour demand at the industry or sector level and the development of labour demand at the aggregate level (Baumol, 1967; Foster et al., 2017). In the case of labour-saving technologies, several compensation mechanisms can counterbalance the initial adverse impact of automation. The countervailing mechanisms that have received most attention in the literature are discussed below, as well as the empirical evidence related to those mechanisms.

4.2.1 Compensation through complementarities

Technological innovations are not only labour-saving, they can also augment the demand for labour. A central economic mechanism through which automation positively affects the demand for labour is that it raises the economic value of tasks that humans uniquely supply. In many cases, machines not only substitute for labour, but they also complement workers. Especially tasks that cannot be substituted by automating technologies are generally complemented by it. The production of goods and services requires combining a continuum of tasks that draw upon a multifaceted set of inputs (equivalently, skills). Typically, these inputs each play essential roles. Consequently, improvements in the productivity of one task or a fall in its price does not reduce the need for other tasks. Hence, if task B is an important, or even an indispensable, complement to task A that is automated, the demand for B will increase. For example, the continuing price decline of computers, in combination with a strong increase in computing power, has led to an increased substitution of computer-assisted technologies for human labour. Routine tasks predominantly involving the organisation, storage, retrieval, and manipulation of information are increasingly codified in computer software and performed by machines. As non-routine abstract tasks are heavily dependent on information as an input, these types of tasks tend to be complemented by computer technologies. By lowering the cost of retrieving and manipulating information, workers in abstract task-intensive jobs will reduce the time spent on routine tasks. Accordingly, computerization enables workers to further specialize in their area of comparative advantage, i.e. analysing and interpreting information. This process of automation, in turn, raises the relative demand for workers who can perform complementary non-routine tasks. As a result, non-routine analytical skills are increasingly rewarded on the labour market (Autor et al., 2003; Handel, 2012; Ingram & Neumann, 2006). As the empirical evidence provided below will illustrate, the strong complementarities between automation and labour can boost productivity, raise earnings, and in turn, increase demand for labour.



4.2.2 Compensation through lower costs

Although technologically advancing sectors typically shrink as a share of employment, the adoption of new technologies can increase employment if the price elasticity of demand is sufficiently elastic (Bessen, 2017). Automation will reduce the prices for tasks and, thereby, the average cost of producing a good or service. In competitive markets, lower production costs will translate into lower prices. If demand increases sufficiently as a response to lower prices, employment may rise even though the labour required per unit of output falls (Acemoglu & Restrepo, 2017). An example of a good or service for which a decrease in price elicits more than a proportional increase in its demand is air travel. Advances in aeronautical engineering led to a reduction in the price of air travel after 1903, which in turn led to an increase in total spending on this type of travel and employment growth in this industry (Brynjolfsson & Mitchell, 2017). Another example of an industry that experienced net job growth despite the adoption of labour-saving technologies is the textile industry (Bessen, 2015; Bessen, 2019). During the 19th century, power looms replaced 98 percent of the labour required to weave a yard of cloth. Nevertheless, the number of weaving jobs rose over the same period. Automation technologies lowered the price for cloth in a competitive market, increasing the highly elastic demand for cloth, which in turn resulted in employment growth (Bessen, 2015). In a similar vein, the primary steel and automotive industries (where demand was initially highly elastic) experienced employment growth at the same time as labour productivity rose (Bessen, 2019). According to Vivarelli (1995, 2014), a price decline is perhaps the most effective compensation mechanism for limiting employment losses (Vivarelli, 1995, 2014). However, in sectors characterized by low price elasticity of demand, the job-destroying impact of technological change might not be offset by compensating market mechanisms. As a result of increases in agricultural labour productivity and falling relative prices of agricultural goods, economy-wide prosperity increases and household demand for agricultural products grew less than demand for other goods. Autor (2014) documents that in 1900, 41 percent of the US workforce was employed in agriculture; by 2000, that share had fallen to 2 percent, mostly due to a wide range of technologies including automated machinery, such as field machinery and irrigation systems. In the case that markets are not perfectly competitive, the decrease in production costs resulting from technological progress will not be fully translated into falling prices. As a result, innovating firms can accumulate extra profits. If these profits are reinvested in the firm, new jobs can be created under the condition that these investments are labour-intensive (Vivarelli, 2014).



4.2.3 Compensation mechanism via increase in incomes

The adoption of technology could raise the productivity of workers and thereby the total income of these workers or the broader population (Acemoglu & Restrepo, 2017). Increases in income are likely to lead to increased consumption of goods and services and an increased derived demand for the tasks needed to produce those goods and services. As such, rising income may also spur demand for activities outside industries or sectors that are subject to technological change. For example, as total income has increased, Americans have spent more of their income on restaurant meals (Brynjolfsson & Mitchel, 2017). This increase in product demand results in a rise in employment which may fully compensate for the initial job losses caused by automation technologies (Boyer, 1988a, 1988b; Pasinetti, 1981). However, if the share of the total income is not distributed equally, income inequalities may also rise sharply.

4.2.4 Empirical evidence

The direct effect of technological progress on employment and labour share in the firms, sectors or industries that are subject to it is generally easily observable. Because the demand for labour tends to contract in technologically advancing sectors (Baumol, 1967), the *direct* labour-displacing effects shape the empirical analyses and public debate of the aggregate (net) impact of technological change. The *indirect* effects of technology adoption are more difficult to observe and quantify, and may consequently receive less attention. Autor and Salomons (2018) find that at the aggregate level, technological change – measured by productivity growth – has been labour-augmenting rather than labour-reducing. Using a sample of 32 industries in 19 developed countries between 1970 and 2007, the authors show that the direct labour-reducing effect in industries in which technological innovations originate, is more than offset by a rise in country-level employment as aggregate productivity grows. This compensating effect is realised through two channels. First, a productivity rise in supplier industries leads to strong employment gains in downstream customer industries. Second, productivity growth raises total income and hence final demand, which in turn induces employment growth across all sectors. Likewise, Bessen (2017) finds that the introduction of computer technology has reduced the demand for labour in manufacturing industries, but mildly increased employment in non-manufacturing industries in the United States between 1984 and 2007. The heterogeneous employment effects can be explained by differences in the price elasticity of demand. The price elasticity of demand in manufacturing industries has become rather low due to ongoing productivity gains and met demand. In non-manufacturing industries that are characterized by high prices and large unmet demand,



productivity improvements and price declines will spur the demand for products and therefore have a more positive impact on employment.

Using data from the International Federation of Robotics (IFR), Graetz and Michaels (2018) examine the employment impact of industrial robot use in 14 industries across 17 countries. Averaged across countries, robot density increased by more than 150% between 1993 and 2007. The authors demonstrate that industrial robots contributed positively to productivity growth, while at the same time reducing output prices. The adoption of robots has had positive wage effects and no adverse aggregate employment effects. The positive effect on wages suggests that some of the productivity gains from increased robot use were shared with workers. Albeit industrial robots did not significantly reduce total employment, increased robot density did reduce the employment share of low-skilled workers. The IFR data is also used by Acemoglu and Restrepo (2017), but their analysis of the employment impact of robots is restricted to the United States. In contrast to Graetz and Michaels (2018), the authors document that regions most exposed to industrial robots between 1990 and 2007 have experienced a decline in wages and employment. Dauth et al. (2017) find no evidence for Germany that robot exposure affects the overall level of local employment, while it does change the employment composition. The decline in manufacturing jobs is fully offset by additional jobs in the service sector. Robot exposure raises the wages of high-skilled workers, but mainly declines for middle-skilled workers. Gregory, Salomons and Zierahn (2016) estimate the impact of routine-tasks replacing technologies for 238 regions across 27 European countries over the period 1999-2010. In aggregate, routine-replace technological change has net resulted in positive labour demand effects. Gregory et al. (2016) show that compensating product demand and local demand spill-over effects dominate the direct labour-displacing effect of technological change. Likewise, Autor et al. (2015) find no evidence for a negative employment effect of exposure to computerization of routine-intensive tasks in local labour markets. The general conclusion from this literature is that technology has had small (and mostly positive) effects on the overall level of employment.

4.2.5 Impact: potential outcomes

From these considerations, we can draw a number of conclusions about plausible future developments. These predictions are assumptions we deem plausible, but not beyond discussion. They do however form the foundation under our future scenarios. First, we may assume that the adoption of key technologies will substitute for workers in some job tasks, and augment workers in others. These developments will plausibly occur simultaneously. This will plausibly lead to



disappearance of some jobs, and the moulding of most other jobs. A number of jobs will also be created, but it is yet unknown how many, and also what skills maybe required. It will be reasonable to assume that the task profile of future jobs will be different from most of today's jobs, and that they will require different skills of workers in these jobs. At the same time, increased productivity will spill over to other firms and sectors, creating employment in these sectors.

For reasons of simplicity, and to maximize the utility of our scenarios for provoking thoughts about potential impact, we consider the two outcomes we consider most extreme. In the first outcome, we assume that intelligent automation will **mostly substitute human labour**, and that a maximum number of jobs will disappear. If we may take the most upper-bound predictions as a guideline, this may amount to about half of all jobs (Frey and Osborne, 2017). The second scenario assumes that intelligent automation will **mostly augment human labour**. In this scenario, the net effect of automation on the number of jobs will be close to zero, which means that most jobs that disappear will be replaced by new jobs, and that existing jobs may change to require new skills.

5 Scenarios for the future of work

From the previous sections, it becomes clear that we believe three variables to be of key concern for shaping the impact of automation on work, i.e. 1) the speed of innovation, 2) the speed of adaptation, and 3) the impact on job tasks. For these three variables, for reasons of simplicity, we describe two potential future outcomes for each of these variables.

By combining the potential outcomes of the three variables, we arrived at eight scenarios for the impact of innovations on future work. We describe these scenarios from the perspective of the years 2025 or 2035, looking back on possible developments between today and then. Using this thought experiment, we describe possible outcomes that would come along with the different “types of future”. These scenarios serve to underline that the future is not fixed, but the way the future unfolds will (partly) depend on these important variables. In drafting these scenarios, we also make a number of assumptions on other important variables, including the following:

- Current demographic trends in most European countries (e.g. population aging, declining fertility rates, net migration rates) can be extrapolated into the future,
- Sociological trends (e.g. educational expansion, increasing social homogeneity, rising social inequalities) continue into the future,
- Governments are not passive actors but are able to respond proactively to expected trends. Following current policy debates, we assume they will mostly do so by reforming (a) labour market policies, (b) social welfare programmes, and (c) education and training systems. Of course, large cross-national differences in policy responses will occur, but we reason that all countries will reform policies along these dimensions.

Of course, these scenarios follow from a stylized and simplified model of reality. The scenarios we describe are neither mutually exclusive nor exhaustive and should not be mistaken for predictions. They describe how the world *could* be, given reasonable assumptions, and not how the world *will* be. It is not expected that the future will unfold precisely following either of these scenarios; rather, these scenarios will probably occur concomitantly, with some scenarios being more plausible in some economic sectors and countries than in others.

Although logic dictates that some of the potential outcomes are similar in alternative but comparable scenarios, we chose to emphasize different potential developments in different outcome scenarios, describing scenarios for the years 2025 and 2035. This allows us to present and



discuss a wider range of potentially interesting policy and business dilemmas that follow from intelligent automation.

Eight possible future scenarios

1. Acute disruption

Innovation: boom

Adoption: fast

Impact on work: mostly substitution

2. Incremental automation

Innovation: boom

Adoption: slow

Impact on work: mostly substitution

3. Delayed disruption

Innovation: gradual

Adoption: fast

Impact on work: mostly substitution

4. Slow substitution

Innovation: gradual

Adoption: slow

Impact on work: mostly substitution

5. Abrupt volatility

Innovation: boom

Adoption: fast

Impact on work: mostly augmenting

6. Controlled adjustment

Innovation: boom

Adoption: slow

Impact on work: mostly augmenting

7. Delayed volatility

Innovation: gradual

Adoption: fast

Impact on work: mostly augmenting

8. Gradual evolution

Innovation: gradual

Adoption: slow

Impact on work: mostly augmenting



5.1 Acute disruption

Innovation: boom

Adoption: fast

Impact on work: mostly substitution

In this scenario, the main engineering bottlenecks of crucial technologies (AI, computing, and/or robotics) were overcome by 2025, and as a result, machines have been able to perform both routine and non-routine job tasks, both in the service sector and in various industrial sectors. More or less concomitant to these developments, companies in important economic sectors jumped on the technological bandwagon, sped up the adoption of key technologies so as to keep up with the competition, succeeded in overcoming the main organisational and cultural roadblocks to technology implementation, and adopted their HR policies to facilitate a transition to implement technologies in their value chains.

As a result, the companies who implemented these technologies created new roles and functions to achieve smooth integration of technologies into work processes, and moulded job contents to include tasks that would allow human workers to maximize productivity by working with the new technologies. Companies initially repositioned and retrained their staff wherever that was required, possible, and cost-efficient, to fill in the new job roles demanded by technological developments. Initially, these roles were geared toward programming and implementing AI and digitalization of internal processes, and generating training data based on a successful re-design. However, as AI's became trained better in more processes, intelligent automation was increasingly able to perform routine and non-routine tasks.

At this point, the demand for workers dropped, and companies started to let off human workers whose acquired skills sets no longer matched the skills required in the new organisation, and for whom repositioning was not an option. This led to large numbers of jobs disappearing and large numbers workers becoming unemployed. Wages of workers with demanded skills rose precipitously, but on average, wages steeply dropped. Inequalities soared. Because this happened in a relatively short period of time, the labour market did not have time to adjust. As a result, large groups of workers became unemployed for a long period of time. Education systems were not yet capable of supplying the labour market with required skills, because it long remained unclear what skills would be in demand in the new economy. By the time new jobs became available, however, new cohorts of school-leavers were educated to be proficient in skills demanded to be productive.



This supply of young, agile, trained and cheap labour further hampered reintegration of the technologically unemployed.

5.2 Incremental automation

Innovation: boom

Adoption: slow

Impact on work: mostly substitution

In this future, the main engineering bottlenecks of AI, computing, and/or robotics were overcome by 2025, and it became clear what the potential of these technologies for intelligent automation of work was. However, market diffusion took a much longer trajectory, because of organisational, regulatory and cultural barriers, and because of practical issues such as lagged process digitalization and a delay in the spread of paperless offices. As a consequence, companies early on prepared for implementation of intelligent automation. Investments in HR strategies were made in time and cultural barriers were eventually overcome by specifically designed policies. As a result, the companies who implemented these technologies were able to mold jobs profiles to the demands of the new technologies and thus created new roles and functions to guarantee smooth integration of technologies into work processes. However, all of that took much time.

As a result, governments had ample time to reform education systems to ensure that children were learning the required skills. The stock of well-skilled school-leavers was too low to meet the initial demand, so reskilling of workers became essential. Companies also invested and retrained staff where needed and cost-efficient for adapted roles. They mainly selected the most talented from their own rank and file for retraining or hired workers that already acquired the occupational-specific skills required, either in formal, non-formal, or informal education. Wages of workers with demanded skills increased considerably, as companies competed for skilled workers. These workers had increasingly interesting jobs that relied heavily on creativity, problem solving, and complex communication.

However, with the demand for human labour steadily declining, companies also incrementally let off human workers whose skills did not match the skills required in the new organisation, and for whom retraining was not an option. This led to a slowly but steadily growing surplus on labour markets, particularly for those with lower learning skills and those whose occupational-specific skills were no longer in demand. The labour market partly adjusted for this, as lower-skilled service jobs increased in numbers and new jobs were created. After some time trying to find work, some



workers thus found new careers, mostly in jobs that intelligent machines cannot do. Over-education increased, and some crowding out of lower educated workers took place. Workers that could not find new jobs eventually became inactive on the labour market. On average, incomes steeply dropped, and more people became dependent on income from other sources than labour. Particularly older workers retired early.

5.3 Delayed disruption

Innovation: gradual

Adoption: fast

Impact on work: mostly substitution

Tech companies struggled to overcome the main engineering bottlenecks of AI, computing, and/or robotics for over a decade. Slowly increasing data availability in many processes hampered training AI for many processes. The slow and gradual innovation pace gave companies ample time to prepare adoption of key technologies and deal with organisational and cultural roadblocks to the implementation of technology in their value chains. Because most new technologies did not yet reach their potential, many companies were implementing tech that could not yet fulfil all the expectations. Some companies got disappointed and lost interest in investing in new technologies. Those that did remain, had overcome the most important organisational and cultural barriers once the main bottlenecks were solved, and were ready to adopt new technologies fast and relatively smoothly. This allowed them to gain significant early competitive advantages over those companies that did not. The following competition exponentially increased the pace of tech adaptation.

All of this had consequences for workers. At first, wages for workers with demanded skills steeply increased as companies competed for tech talent. By the time intelligent automation technologies reached their full potential, the early adopters had geared their organisations towards smooth integration of intelligent automation technologies that were already able to perform many routine and non-routine tasks. The demand for workers with skills that machines were also proficient in, steeply dropped in a short period of time. Organisations started to save on human wage costs, either by letting off redundant personnel or by suppressing wages. Unemployment peaked and quickly turned into long-term inactivity. Efficient production meant lower prices. Other companies had to quickly follow suit or be driven out of business.

The relatively slow pace of tech diffusion meant that governments would have had sufficient time to understand the capabilities of technologies and prepare their education systems for the future. New cohorts of school-leavers were more likely equipped with skills that were in demand on the labour market. However, with less work to be done for humans, non-standard work forms (temporary work, part-time work, mini-jobs) became increasingly common for many. Increasing numbers of people, both workers, unemployed, and inactive, became dependent on support income from other sources than work.



5.4 Slow substitution

Innovation: gradual

Adoption: slow

Impact on work: mostly substitution

In this scenario, innovation of key technologies took at least until 2035, and the diffusion of these technologies was also very slowly paced. By the time intelligent automation reached its full potential and diffusion in the labour market reached saturation, the labour markets and societies of most European countries had fundamentally changed. Most notably, many older workers had retired and were no longer active on the labour market. As such, the supply of skilled labour had already significantly decreased and the ratio of workers to retirees had gradually dropped in many European countries.

The pressure on social support systems was exacerbated by the consequences of slowly but steadily advancing automation for younger, lower skilled workers. The slow pace of innovation meant that companies had considerable time to prepare their organisations for the implementation of intelligent automation. Job moulding and limited reskilling ensured that the most skilled personnel remained employed for a long time. As the demand for human labour gradually declined with increasing machine capabilities, companies had incrementally let go of human workers with obsolete skills. These workers mostly did not find new employment, and increasingly relied on social support for their livelihood. Pension systems and social welfare systems were increasingly costly, and the average income of people dependent on support, dropped.

Education systems also where gradually reformed, and labour market arrangements had been put in place to ensure labour markets could adjust gradually.

5.5 Abrupt volatility

Innovation: boom

Adoption: fast

Impact on work: mostly augmenting

The innovation of technologies crucial to intelligent automation progressed fast and reached its zenith around 2025, when most important engineering bottlenecks had been solved. Since that moment, intelligent machines had been capable of performing both routine and non-routine job tasks in many economic sectors. Companies had rapidly adopted these technologies and were able to overcome the most important organisational and cultural bottlenecks as well. Intelligent automation technologies were rapidly integrated in companies' value chains on a large scale.

Some jobs disappeared, but most jobs were moulded in such a way that the potential productivity gains from intelligent automation were maximized. Consequentially, the demand for workers with skills that were complementary to machines increased rapidly. However, these skills were not readily in supply, as many workers were educated for jobs that did not yet rely on intelligent automation.

In response, four developments ensued. First, companies competed heavily for workers with the desired skills, whose wages increased rapidly. Second, international and intra-national migration of skilled workers rose sharply, as a response to this competitiveness. Third, large numbers of workers found that their skills had become obsolete and engaged in reskilling. The most talented workers followed retraining programmes with their employers, but many others lost their jobs. Those for whom retraining was an option relied on government-sponsored training programmes or invested in their own training. Others found employment in lower-paid jobs, and others still became inactive.

Fourth, in response to changing skill demands, governments were prompted to reform their state education systems to ensure that school-leavers had the skills required to be productive. In order to ensure competitiveness in the global market, governments rolled out study programmes that would enable school-leavers to contribute to technological innovation. In vocationally oriented education systems, programmes that would allow school-leavers to effectively work with intelligent automation, or endow them with skills that machines did not yet have, were increasingly implemented. In general education systems, the emphasis would remain on teaching active learning techniques that would allow school-leavers to be flexible.

5.6 Controlled adjustment

Innovation: boom

Adoption: slow

Impact on work: mostly augmenting

In 2025, most important engineering bottlenecks for AI and robotics had been overcome and in principle, machines could perform both routine and non-routine job tasks in many economic sectors. However, companies had difficulties overcoming the organisational and cultural bottlenecks that hampered integration of these technologies into their value chains. As a result, the implementation of intelligent automation took a long time.

This gave governments enough time to work with employers and prepare education systems to meet the skill demands required. Many had done so. State education systems were delivering cohorts of skilled school-leavers to the labour market to meet the demand. As a result, macro inefficiencies in the labour market for school-leavers were limited.

The long diffusion time of intelligent automation technology also gave companies the opportunity to gradually invest in human capital required to implement intelligent automation. Investment in human capital became essential for competitiveness. Companies remoulded work and production chains and jobs and trained and hired staff for these adapted roles. They incrementally let go of workers whose skills were obsolete and who could not efficiently be retrained.

Middle income jobs required upskilling or disappeared, and the number of jobs requiring high-level skills increased. Lower income jobs requiring craftsmanship and personal socioemotional intelligence were too expensive to automate and also did not decline in number. To make the transition to a new job, informal education and adult education were essential for workers' productivity and long-term employability. Workers who could not have access to retraining or who could not retrain successfully for upskilling, took jobs below their level of formal education. The extent to which this was the case, depended largely on access to adult education. In some countries, the responsibility for retraining was regarded a responsibility of employers. Here, relatively highly skilled staff more likely had access to retraining, which increased productivity but also inequality. In other countries, adult informal learning was sponsored by state programmes. Here, inequalities were smaller.

5.7 Delayed substitution

Innovation: gradual

Adoption: fast

Impact on work: mostly augmenting

In this scenario, innovation of technologies crucial for intelligent automation took at least until 2035. Bottlenecks were only slowly and gradually overcome. By the time they were marketed, companies were ready to adopt technology and integrate them into their value chains. They had incrementally hired and trained staff to help adopt innovations. Few of these were school-leavers, but as the slow and gradual innovation rate made it impossible for governments to sensibly reform education systems, it was long unclear what the potential of new technologies would be, and thus what skills would be in demand. The gradual innovation rate also meant that the need for educational reform was obscured. In many countries, a sense of urgency for such reforms was missing.

As a result, when intelligent automation was reaching its peak, many workers still had not been retrained for productivity in the new economy. Temporarily, both the number of vacant jobs and unemployment rose. Many workers actively engaged to find new jobs. The labour market on which workers had to operate had gradually polarized, with the number of middle-income jobs decreasing and both lower and higher income jobs surging. Upskilling programmes became increasingly important for gaining high-income work, for those who were able to do so. People in high-income jobs saw their wages increase.

Those who could not participate in training programmes that would upgrade their skills, had to accept lower-income jobs. As a result, over-education became much more prevalent, and downward social mobility became more common. Wages for low-income jobs did not increase markedly, since higher wages would make these jobs likely candidates for automation. As a result, social inequalities increased. Many new jobs in the low-skilled service sector were created. The increased productivity and availability of workers in high-income jobs resulted in longer work weeks, which in turn implied that more people spent less time on household activities. Their increased income also implied that they could hire personnel to perform household work.

5.8 Gradual substitution

Innovation: gradual

Adoption: slow

Impact on work: mostly augmenting

Innovation of key technologies did not really reach its peak before 2035; the diffusion of intelligent automation technologies also happened incrementally. Eventually, intelligent automation reached its potential, and market adoption picked up speed. At that time, population aging and low birth rates had changed the European population. Mass retirement significantly lowered the worker-to-retiree-ratio and the supply of skilled labour had decreased in many European countries.

As the increasing machine capabilities meant that machines had become more and more able to perform routine and non-routine jobs, the demand for skills had gradually changed. Keeping pace with these developments, companies had replaced workers whose skills were no longer in demand with workers who had the required skills on a piecemeal basis. Retraining programmes helped workers to make the transition to a new job. Formal education reforms ensured that school-leavers were taught the skills needed on the labour market.

6 Discussion

With this report we aimed to better understand the potential consequences of technological change and describe possible and plausible future scenarios of how technological innovations may affect labour markets in the EU. These scenario studies aimed to analyse potential future scenarios, by determining what variables would be crucial for the impact of intelligent automation on work, and then reason what would likely happen given the value of these variables. We described eight possible scenarios from the perspective of the years 2025 and 2035, thereby identifying the different developments that may occur in the near future.

In almost all scenarios, the consequences for human work are considerable. Under the assumption that innovations would primarily have a labour-saving effect, reduced demand for labour and rising unemployment (either short term or long-term) is a distinct possibility. Even if technology augments human productivity, it seems a rise in frictional unemployment is likely. In such scenarios, reskilling of the labour force is essential. In all scenarios, labour market efficiency is only achieved when education systems succeed in endowing school-leavers with the skills demanded on the labour market.

The strength of these scenario studies is that it allows some insights for how to think in a structured way about the future, in cases where extrapolating existing trends is not feasible. The model creates a simplified version of reality, that produces clear, distinguishable potential scenarios. Our comprehensive description of how the interaction of crucial variables determine the impact of technological change can help those attempting to quantify the impact of automation on employment (see, e.g., Heald, Smith & Fouarge, 2019).

One weakness of the scenario study design is the inherent simplicity of its underlying assumptions. Although many variables impact the complex future outcomes, in our reasoning we must limit ourselves to the variables that are part of our model. And these assumptions are formulated as simple as possible, to make the various scenarios relatable and understandable. However, if more complex assumptions would have been made, different scenarios would have become entirely plausible. Case in point, in our review of the potential impact of technology on work, our overview of the relation between technological innovations and macro-economic outcomes is not comprehensive. We assume that investments will by definition lead to labour productivity increases. But from a macroeconomic perspective, it may be safe to assume that the link from rising productivity to jobs depends on who captures the economic rent of increasing productivity:



- Consumers, through the price mechanism
- Workers, as the labour market adjusts so that wages reflect labour productivity
- Owners of capital, through increasing profits

If consumers and/or workers benefit, then they will be better off in real terms and it may be safe to assume that they spend this additional income generating jobs elsewhere. However, net job creation for every Euro spent may well decline – possibly radically depending on what new goods and services are bought with this additional income. This may actually be the one difference of the technological revolution compared to earlier ones: the new (additional) goods and services demanded are also produced without any labour input. If economic rents go to capital owners, then the issue is whether this money is re-invested in the real economy or simply stored. The accumulation of large wealth not only has inequality implications but could also mean that the demand side of the economy stalls and so no new jobs are generated.

Such considerations are not part of our models, and readers should keep an open mind toward them.

References

- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter, & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4, pp. 1043-1171). Elsevier.
- Acemoglu, D., & Restrepo, P. (2017). *Robots and jobs: Evidence from US labor markets*. NBER Working Paper No. 23285.
- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488-1542.
- Amazon (2019). *What robots do (and don't do) at Amazon fulfillment centers*. Retrieved from: <https://www.aboutamazon.com/amazon-fulfillment/our-innovation/what-robots-do-and-dont-do-at-amazon-fulfillment-centers/>
- Ancarani, A., Di Mauro, C., Legenvre, H., & Cardella, M. S. (2019). Internet of things adoption: A typology of projects. *International Journal of Operations & Production Management*, (forthcoming).
- Arntz, M. T., Gregory, T., & Zierahn, U. (2016). *The risk of automation for jobs in OECD countries: A comparative analysis*. OECD Social, Employment and Migration Working Papers, No. 189, OECD Publishing, Paris.
- Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G. S. L., Buell, D. A., Burkett, B., Chen, Y., Chen, Z., Chiaro, B., Collins, R., Courtney, W., Dunsworth, A., Farhi, E., Foxen, B., Fowler, A., Gidney, C., Giustina, M., Graff, R., Guerin, K., Habegger, S., Harrigan, M. P., Hartmann, M. J., Ho, A., Hoffmann, M., Huang, T., Humble, T. S., Humble, T. S., Isakov, S. V., Jeffrey E., Jiang, Z., Kafri, D., Kechedzhi, K., Kelly, J., Klimov, P. V., Knysh, S., Korotkov, A., Kostrița, F., Landhuis, D., Lindmark, M., Lucero, E., Lyakh, D., Mandrà, S., McClean, J. R., McEwen, M., Megrant, A., Mi, X., Michielsen, K., Mohseni, M., Mutus, J., Naaman, O., Neeley, M., Neill, C., Yuezhen Niu, M., Ostby, E., Petukhov, A., Platt, J.C., Quintana, C., Rieffel, E. G., Roushan, P., Rubin, N. C., Sank, D., Satzinger, K. J., Smelyanskiy, V., Sung, K. J., Trevithick, M. D., Vainsencher, A., Villalonga, B., White, T., Yao, J., Yeh, P., Zalcman, A., Neven, H., & Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505-510.
- Autor, D. H. (2014). *Polanyi's paradox and the shape of employment growth*. NBER Working Paper No. 20485.



- Autor, D. H. (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. H., Dorn, D., & Hanson, G. H. (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125(584), 621-646.
- Autor, D. H., & Salomons, A. (2018). *Is automation labor-displacing? Productivity growth, employment, and the labor share*. NBER Working Paper No. 24871.
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279-1333.
- Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of urban crisis. *The American Economic Review*, 57(3), 415-426.
- Beal, G.M., Rogers, E.M., Bohlen, J.M. (1957). Validity of the concept of stages in the adoption process. *Rural Sociology*, 22(2), 166–168.
- Benzell, S. G., Kotlikoff, L. J., LaGarda, G., & Sachs, J. D. (2015). *Robots are us: Some economics of human replacement*. NBER Working Paper No. 20941.
- Berry, F. S., & Berry, W. D. (2018). Innovation and diffusion models in policy research. In C. M. Weible, & P. A. Sabatier (Eds.), *Theories of the policy process* (pp. 263-308). Routledge.
- Bessen, J. (2015). Toil and technology: Innovative technology is displacing workers to new jobs rather than replacing them entirely. *IMF Finance and Development Magazine*, 52(1), 1-4.
- Bessen, J. (2017). *Automation and job: When technology boosts employment*. Boston University of Law, Law and Economics Research Paper No. 17-09.
- Bessen, J. (2018). *AI and Jobs: The role of demand*. NBER Working Paper No. 24235.
- Bessen, J. (2019). Artificial intelligence and jobs: The role of demand. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 291-307). University of Chicago Press.
- Bessen, J. E., Goos, M., Salomons, A., & Van den Berge, W. (2019). *Automatic reaction-What happens to workers at firms that automate?* Boston University School of Law, Law and Economics Series Paper No. 19-2.
- Bohlen, J.M. & Beal, G. M. (1957). *The Diffusion Process*. Special Report No. 18. 1: 56–77.
- Boyer, R. (1988a). New technologies and employment in the 1980s: From science and technology to macroeconomic modelling. In J. A. Kregel, E. Matzner, & A. Roncaglia (Eds.), *Barriers to full employment* (pp. 233-272). Palgrave Macmillan, London.
- Boyer, R. (1988b). Technical change and the theory of regulation. In G. Dosi, C. Freeman, R. Nelson, G. Silverberg & L. Soute (Eds.), *Technical Change and the Economic Theory* (pp. 67-94). London, UK: Pinter.

- Brown, N., & Sandholm, T. (2018). Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374), 418-424.
- Brynjolfsson, E. & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530-1534.
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy?. *AEA Papers and Proceedings*, 108, 43-47.
- Burns, T., & Stalker, G. M. (1994). *The management of innovation*. New York, NY: Oxford University Press.
- Chui, M., J. Manyika, & Miremadi, M. (2018, January). *What AI can and can't do (yet) for your business*. McKinsey Quarterly.
- Columbus, L. (2019, September 25). What's New In Gartner's Hype Cycle For AI, 2019. *Forbes*.
- Copeland, M. (2016). *What's the Difference Between Artificial Intelligence, Machine Learning, and Deep Learning?* <https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/> (last access 21.12.2018).
- Dauth, W., Findeisen, S., Südekum, J., & Woessner, N. (2017). *German robots-the impact of industrial robots on workers*. CEPR Discussion Paper No. DP12306.
- DiVincenzo, D. P. (1995). Quantum computation. *Science*, 270(5234), 255-261.
- European Commission (2018). *Communication from the Commission to the European Parliament, the European Council, the Council, The European Economic and Social Committee and the Committee of the Regions: Artificial Intelligence for Europe*. Brussels, 25.4.2018 COM(2018) 237 final. Online: <https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe> (last access 01.10.2019)
- Ford, M. (2015). *Rise of the robots: Technology and the threat of a jobless future*. Basic Books.
- Foster, L. S., Grim, C. A., Haltiwanger, J., & Wolf, Z. (2017). *Macro and micro dynamics of productivity: From devilish details to insights*. NBER Working Paper No. 23666.
- Fountaine, T., McCarthy, B., & Saleh, T. (2019). Building the AI-Powered organization. *Harvard Business Review*, 63-73.
- Fraunhofer-Gesellschaft (Ed.) (2018). *Maschinelles Lernen. Eine Analyse zu Kompetenzen, Forschung und Anwendung [Machine Learning. Analysing competencies, research and deployment]*, München: Fraunhofer.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254-280.



- Gartner. (2019). *Gartner Hype Cycle. Interpreting technology hype*. Retrieved from: <https://www.gartner.com/en/research/methodologies/gartner-hype-cycle>
- Gerbert, P., Justus, J., & Hecker, M. (2017). *Competing in the age of Artificial Intelligence*. The Boston Consulting Group.
- Gesing, B., Peterson, S. J., & Michelsen, D. (2018). *Artificial intelligence in logistics*. Troisdorf, Germany: DHL Customer Solutions & Innovation.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. Cambridge, MA: MIT Press.
- Goolsbee, A. (2018). *Public policy in AI economy*. NBER Working Paper No. 24653.
- Graetz, G., & Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5), 753-768.
- Greenhalgh, T., Robert, G., Macfarlane, F., Bate, P., & Kyriakidou, O. (2004). Diffusion of innovations in service organizations: Systematic review and recommendations. *The Milbank Quarterly*, 82(4), 581-629.
- Heald, S., Smith, A., & Fouarge, D. (2019). Labour market forecasting scenario's for automation risks: Approach and outcomes. Technequality Working Paper (Deliverable D1.4).
- High-Level Expert Group on Artificial Intelligence (AI-HLEG) (2019). *A Definition of AI: Main Capabilities and Disciplines*. Brussels: European Commission. Online: <https://ec.europa.eu/digital-single-market/en/news/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines> (last access: 08.10.2019)
- Gregory, T., Salomons, A., & Zierahn, U. (2016). *Racing with or against the machine? Evidence from Europe*. ZEW – Centre for European Economic Research Discussion Paper No. 16-053.
- Habakkuk, H. J. (1962). *American and British technology in the nineteenth century: The search for labour saving inventions*. Cambridge University Press.
- Handel, M. (2012). *Trends in job skill demands in OECD countries*. OECD Social, Employment and Migration Working Papers No. 143, OECD Publishing, Paris.
- Harrison, R., Jaumandreu, J., Mairesse, J., & Peters, B. (2014). Does innovation stimulate employment? A firm-level analysis using comparable micro-data from four European countries. *International Journal of Industrial Organization*, 35, 29-43.
- Hawksworth, J., Berriman, R., & Goel, S. (2018). *Will robots really steal our jobs? An international analysis of the potential long term impact of automation*. PricewaterhouseCoopers.
- Hilbert, M., & López, P. (2011). The world's technological capacity to store, communicate, and compute information. *Science*, 332(6025), 60-65.

- Hildesheim, W. (2019) *Artificial Intelligence. What works? What does not work? Growth for and with us. Opportunities for new jobs in Europe*. Presentation at the Technequality Expert Meeting, April 24-25, 2019.
- Hogarth, T. (2017). *Automation artificial intelligence on-demand labour and other flexible forms of labour in the new IDB employer survey Skills at Work in LAC*. Inter-American Development Bank Technical Note No. IDB-TN-1332.
- Ingram, B. F., & Neumann, G. R. (2006). The returns to skill. *Labour Economics*, 13(1), 35-59.
- International Federation of Robotics (2017). World Robotics 2016 Industrial Robots. Frankfurt am Main, Germany: IFR.
- Keynes, J. M. (2010). Economic possibilities for our grandchildren. In *Essays in persuasion* (pp. 321-332). Palgrave Macmillan, London.
- Lefranc, A. (2003). *Labor market dynamics and wage losses of displaced workers in France and the United States*. William Davidson Institute Working Paper No. 614.
- Manyika, J., Lund, S., Chui, M., Bughin, J., Woetzel, J., Batra, P., & Sanghvi, S. (2017, November). *What the future of work will mean for jobs, skills, and wages*. McKinsey Quarterly.
- McKinsey (2018, November). *AI adoption advances, but foundational barriers remain*. New York: McKinsey.
- McKinsey Global Institute (MGI; 2018). *Notes from the AI-frontier (Discussion paper)* <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning?cid=other-eml-ttn-mip-mck-oth-1806&hlkid=d94b55c44cfb4258b7dc49f328629182&hctky=2491952&hdpid=81ba10a7-367d-45eb-96ab-f35ebad887ec>, last acc. 2018-08-10
- Mercer. (2019). *Global Talent Trends Study – The Future of Work*. New York: Mercer.
- Miller, M. (2019, October). Gartner: The Present and Future of Artificial Intelligence. *PCMag* <https://www.pcmag.com/article/371612/gartner-the-present-and-future-of-artificial-intelligence>
- Mintzberg, H. (1979). *The structuring of organizations*. London, England: Pearson.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533.
- Mohanty, S., & Vyas, S. (2018). *How to compete in the age of artificial intelligence. Implementing a collaborative human-machine strategy for your business*. New York: Apress.

- Moravčík, M., Schmid, M., Burch, N., Lisý, V., Morrill, D., Bard, N., Davis, T., Waugh, K., Johanson, M., & Bowling, M. (2017). Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337), 508-513.
- Moravec, H. (1988). *Mind Children*. Harvard University Press
- Nedelkoska, L., & Quintini, G. (2018). *Automation, skills use and training*. OECD Social, Employment and Migration Working Papers, No. 202. OECD Publishing, Paris.
- Ng, A. (2016, November 9). *What artificial intelligence can and can't do right now*. Retrieved from: <https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now>
- Pasinetti, L. (1981). *Structural change and economic growth*. Cambridge: Cambridge University Press.
- Peters, A. (2019). *Cobots*. Presentation at the Technequality Expert Meeting, April 24-25, 2019.
- Pew Research Centre (2017). *Automation in everyday life*. Retrieved from <https://www.pewresearch.org/internet/2017/10/04/automation-in-everyday-life/>
- Raposo, P. S., Portugal, P., & Carneiro, A. J. M. (2015). *Decomposing the wage losses of displaced workers: the role of the reallocation of workers into firms and job titles*. IZA Discussion Paper No. 9220.
- Salian, I. (2018, August 2). SuperVize Me: What's the difference between supervised, unsupervised, semi-supervised and reinforcement learning? Retrieved from <https://blogs.nvidia.com/blog/2018/08/02/supervised-unsupervised-learning/>
- Scholtes, J. (2019) *Robots in the open*. Presentation at the Technequality Expert Meeting, April 24-25, 2019.
- Shaikh, A. A., & Karjaluoto, H. (2015). Mobile banking adoption: A literature review. *Telematics and Informatics*, 32(1), 129-142.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., Van den Driessche, G., Graepel, T., & Hassabis, D. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676), 354-359.
- Susskind, R. E. (2017). *Tomorrow's lawyers: An introduction to your future*. Oxford University Press.
- Teece, D. J. (1996). Firm organization, industrial structure, and technological innovation. *Journal of Economic Behavior & Organization*, 31(2), 193-224.



- Tractica (2019). *Artificial Intelligence Market Forecasts*. Boulder, CO: Tractica.
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS quarterly*, 27(3), 425-478.
- Vivarelli, M. (1995). The economics of technology and employment. *Edward Elgar Publishing*.
- Vivarelli, M. (2014). Innovation, employment and skills in advanced and developing countries: A survey of economic literature. *Journal of Economic Issues*, 48(1), 123-154.
- Waldrop, M. M. (2016). The chips are down for Moore's law. *Nature News*, 530(7589), 144.
- Williamson, O. E. (1975). Markets and hierarchies. In D. Faulkner (Ed.), *Strategy: Critical perspectives on business and management* (Vol. 4, pp. 106-118). London, England and New York, NY: Routledge.
- World Bank (2019). *World development report 2019: The changing nature of work*. Washington, DC: World Bank.
- Yang, G., Bellingham, J., Dupont, P. E., Fischer, P., Floridi, L., Full, R., Jacobstein, N., Kumar, V., McNutt, M., Merrifield, R., Nelson, B. J., Scassellati, B., Taddeo, M., Taylor, R., Veloso, M., Wang, Z. L., & Wood, R. (2018). The grand challenges of Science Robotics. *Science Robotics*, 3(14), eaar7650.

APPENDIX

Technequality Expert Meeting

Date: 24th & 25th of April 2019

Location: De Nieuwe Liefde, Amsterdam

Goal Expert Meeting

The Technequality consortium organised an Expert Meeting that brought together three types of experts. First, experts at the forefront of creating disruptive technologies provided insight into new job opportunities and into the type of job tasks that can be automated. Second, experts with relevant knowledge of skill requirements in occupations and labour markets shed light on plausible scenarios for the adoption of new technologies and the implications for the future of work. Third, policy experts from various levels of government formulated policy implications based on the discussions in the Expert Meeting. This meeting provided direct input for the innovative models our consortium is building for assessing implications of technological innovations for European labour markets, (2) helped us to better understand how innovations are most likely to change European economies, and (3) helped our scholars to co-create policies with European governments and the European Commission.

Presenting participants and topics presented:

- Glenda Quintini (Senior Economist at OECD): “How automation is affecting skill demand and what we can do about it”
- Jesse Scholtes (Project Manager Robotics at Eindhoven University of Technology): “Robots in the open – development challenges from an engineering perspective”
- Arwen Peters (Marketing & Communication Coordinator at Olmia Robotics): “Are cobots changing the way we work?”
- Pascale LeBlanc (Associate Professor at Eindhoven University of Technology) & Hannah Berkers (Postdoctoral Researcher at Eindhoven University of Technology): “Working with or against the machine: how to optimize human-robot collaboration in the workplace”
- Dr. Wolfgang Hildesheim (Director of Watson & Artificial Intelligence, IBM DACH): “AI – What works? What does not work? Opportunities for new jobs in Europe”
- Prof. Michel Dumontier (Professor of Data Science at Maastricht University): “Embracing the emerging AI economy with FAIR data and services”



- Dr. Carl Frey (Co-Director of the Oxford Martin Programme on Technology and Employment at the University of Oxford): “Automation and work: past, present and future”
- Prof. Steven Dhondt (Senior Researcher at TNO): “Taming the potential. Understanding the competences of companies to deal with the future”
- Tim Schokker & Ted Reininga (Dutch Ministry of Education, Culture and Science): “Netherlands 2040: education and future labour market”
- Ieva Reine (Swedish Social Insurance Agency), Joanna Napierala (European Commission), Santo Milasi (European Commission): Panel discussion

Non-presenting participants

Michiel Bal (KU Leuven); Maaïke Bierman (Maastricht University); Dr. Martin Ehlert (Berlin Social Science Center); Prof. Didier Fouarge (Maastricht University); Dr. Marie-Christine Fregin (Maastricht University); Prof. Andries de Grip (Maastricht University); Sophie Heald (Cambridge Econometrics); Dr. Pantelis Koutroumpis (University of Oxford); Siim Krusell (Kutsekoda); Prof. Mark Levels (Maastricht University); Dr. Raymond Montizaan (Maastricht University); Annarosa Pesole (European Commission); Sandra Reuse (Advisor German Federal Ministry of Labour and Social Affairs); Dr. Triin Roosalu (Tallinn University); Dr. Eve-Liis Roosmaa (Tallinn University); Prof. Ellu Saar (Tallinn University); Sven Semet (IBM); Alistair Smith (Cambridge Econometrics); Prof. Heike Solga (Berlin Social Science Center); Melline Somers (Maastricht University); Prof. Rolf van der Velden (Maastricht University); Dr. Robert Went (Wetenschappelijke Raad voor het Regeringsbeleid)

