



European
University
Institute

ROBERT
SCHUMAN
CENTRE FOR
ADVANCED
STUDIES

WORKING PAPERS

RSCAS 2019/58
Robert Schuman Centre for Advanced Studies
Florence School of Regulation

The economics of Artificial Intelligence: A survey

Laura Abrardi, Carlo Cambini, Laura Rondi

European University Institute

Robert Schuman Centre for Advanced Studies

Florence School of Regulation

The economics of Artificial Intelligence: A survey

Laura Abrardi, Carlo Cambini, Laura Rondi

EUI Working Paper **RSCAS** 2019/58

This text may be downloaded only for personal research purposes. Additional reproduction for other purposes, whether in hard copies or electronically, requires the consent of the author(s), editor(s). If cited or quoted, reference should be made to the full name of the author(s), editor(s), the title, the working paper, or other series, the year and the publisher.

ISSN 1028-3625

© Laura Abrardi, Carlo Cambini, Laura Rondi, 2019

Printed in Italy, July 2019

European University Institute

Badia Fiesolana

I – 50014 San Domenico di Fiesole (FI)

Italy

www.eui.eu/RSCAS/Publications/

www.eui.eu

cadmus.eui.eu

Robert Schuman Centre for Advanced Studies

The Robert Schuman Centre for Advanced Studies, created in 1992 and currently directed by Professor Brigid Laffan, aims to develop inter-disciplinary and comparative research on the major issues facing the process of European integration, European societies and Europe's place in 21st century global politics.

The Centre is home to a large post-doctoral programme and hosts major research programmes, projects and data sets, in addition to a range of working groups and *ad hoc* initiatives. The research agenda is organised around a set of core themes and is continuously evolving, reflecting the changing agenda of European integration, the expanding membership of the European Union, developments in Europe's neighbourhood and the wider world.

For more information: <http://eui.eu/rscas>

The EUI and the RSCAS are not responsible for the opinion expressed by the author(s).

Florence School of Regulation

The Florence School of Regulation (FSR) is a partnership between the Robert Schuman Centre for Advanced Studies (RSCAS) at the European University Institute (EUI), the Council of the European Energy Regulators (CEER) and the Independent Regulators Group (IRG). Moreover, as part of the EUI, the FSR works closely with the European Commission.

The objectives of the FSR are to promote informed discussions on key policy issues, through workshops and seminars, to provide state-of-the-art training for practitioners (from European Commission, National Regulators and private companies), to produce analytical and empirical researches about regulated sectors, to network, and to exchange documents and ideas.

At present, its scope is focused on the regulation of Energy (electricity and gas markets), Communications & Media, and Transport.

This series of working papers aims at disseminating the work of scholars and practitioners on current regulatory issues.

For further information

Florence School of Regulation
Robert Schuman Centre for Advanced Studies
European University Institute
Casale, Via Boccaccio, 121
I-50133 Florence, Italy
Tel: +39 055 4685 878
E-mail: FSR.Secretariat@eui.eu
Web: <http://fsr.eui.eu/>

Abstract

The current advances in Artificial Intelligence are likely to have profound economic implications and bring about new trade-offs, thereby posing new challenges from a policymaking point of view. What is the socio-economic impact of these new technologies on growth, employment and inequality? How markets and competition will be affected by AI-powered agents? What are the implications in terms of consumers' privacy? Will algorithms reduce consumers' biases or will they rather originate new ones? This work presents a first attempt to provide a comprehensive survey of the growing economic literature on Artificial Intelligence and its economic impact on markets and society, focusing on those issues where AI is likely to pose the most imminent challenges for policymakers.

Keywords

Artificial Intelligence, machine learning, algorithms

JEL codes: D24, E24, L50, L86, O33

1 Introduction

The Oxford English Dictionary broadly defines artificial intelligence (AI) as “the theory and development of computer systems able to perform tasks normally requiring human intelligence”. And indeed, computer systems can now perform tasks such as understanding natural language, diagnosing diseases, making up jokes and even driving a car. But the intelligence of computer now goes well beyond their ability to understand and analyze. They can learn. Through Machine Learning (ML), a branch of computational statistics, computers can produce new knowledge by finding complex structures and patterns in example data.

So far, most of the work in the field of AI and ML has had largely an engineering and computer-science orientation. However, quite importantly, AI configures itself as a general purpose technology (Bresnahan and Trajtenberg, 1995; Cockburn et al., 2018), and as such it can be applied transversally across sectors. Indeed, nearly every industry is currently dealing with a surge of intelligent computer systems employed effectively as a decision-aiding tool. Some researchers even arrived at the definition of “algorithmic business” (Ezrachi and Stucke, 2016), i.e. the use of complex algorithms to improve business decisions and automatize processes for competitive differentiation.

Precisely because of their varied and multi-faceted applications, data-driven innovations are having extensive effects on a number of economic levels. Most of the economic literature on AI technologies has focused on their impacts on growth and employment. However, they are also having massive effects on the functioning of existing markets, shifting the mode and efficiency of competition and raising the attention of antitrust and privacy authorities. We provide here a survey of the most recent economic literature, focusing on those aspects of new technologies that pose the most imminent challenges from a policymaking point of view. To this aim, we describe both the effects of the new technologies on productivity and employment, and its implications in terms of competition, collusion and privacy. From our review it emerges clearly that, while the literature on the impacts of new technologies on the labour market is well developed, and firmly connected with that studying the effects of automation on productivity, the impact of machine learning and AI on markets, competition and consumers’ surplus are the object of a rather fragmented and still nascent literature,

despite its relevance.

Under a macroeconomic lens, AI could certainly provide the solution to the problem of stagnant productivity, which has been afflicting the most developed economies for some years now. Indeed, smart machines are a powerful capital-labor hybrid, being able to produce output at scale, but at the same time requiring data more than money, and being capable of learning from data and act accordingly. These peculiar features are generating some degree of dispute within the academic community about how AI can be conceptualized. So far, most models of AI formalize it as automation, i.e. a capital-augmenting factor which substitutes the less productive labour (Kotlikoff and Sachs, 2012; Graetz and Michaels, 2015; Nordhaus, 2015); there is also someone who looks at it as a labor-augmenting factor on the grounds of its complementarity with labour (Bessen, 2018). However, recent work is starting to depart from the automation framework and regards AI as a completely new input of production – for example, research activities, which discover new tasks (Acemoglu and Restrepo, 2018a; Aghion et al., 2018), robots, working in association with labour and capital (DeCanio, 2016), or new decision-making tools for the support of commercial practices (Calvano et al., 2018; Athey et al., 2018).

Regardless of how AI is cast within the productivity framework, there is substantial agreement that it shall have a relevant impact on employment and welfare: just as mechanical muscles made human labor less in demand, so will mechanical minds make human brain less in demand. However, what exactly this impact will be is, once again, an open debate, and conflicting philosophies have been proposed. According to the most futuristic-looking scholars, it will be possible to build a general AI capable of out-performing human labor under any aspect. Then, society will have to deal with what has been defined economic singularity (Nordhaus, 2015): an economy of radical abundance characterized by unbounded growth, in which no one will need to work anymore. In this case, the policymakers' concern should be directed at designing efficient ways to distribute wealth and eliminating market imperfections, so that anyone will benefit from the wealth produced by an unreachable super-intelligence (Korinek and Stiglitz, 2017). A more balanced view argues that, to the extent that AI will replace humans in routine and repetitive jobs, the issues of inequality and unemployment will surge center-stage in the political discussion (Agrawal, Gans and

Goldfarb, 2019). Solving them will require a concerted action of redistribution of wealth, and will possibly entail to devise ways to train people to work on other tasks (Korinek and Stiglitz, 2017). Finally, a more skeptical scenario contends that humans will still be better at thinking outside the box for many years to come (Boden, 1998), and thus AI will mostly be used for work augmentation, providing humans insights, advice and guidance to increase the firm's productivity.

Although the repercussions of AI and ML on productivity and employment have rightly dominated the economic discussion so far, AI will plausibly have far-reaching consequences also at a more microeconomic level, challenging the current framework of the market mechanisms and of the consumers' decision-making processes. The digital economy is characterized by an enhanced role of data, by very strong economies of scale and of scope and by extreme network effects. These key features give rise to a strong incumbency advantage, and originates highly concentrated markets with few dominant players (Cremer et al., 2019). Understanding such breakthroughs and their effects in terms of competition policy, privacy and the efficient allocation of resources (including data) will be paramount for exploiting the benefits and addressing the threats of the new technologies. If, on the one hand, consumers' choices become more accessible, practical and efficient, on the other hand the over-pervasiveness of new technologies could have detrimental side-effects. First of all, consumers might be harmed depending on how and to what extent the information is filtered and used (Calvano et al., 2019). Secondly, the use of consumers' personal information raises significant privacy concerns (Acquisti et al., 2016; Tucker, 2018). Third, as people become more and more reliant on machines for making all the important decisions, machines will take the effective control, with people dependent on them and afraid to make their own choices (Makridakis, 2017).

Consumers' harm could also have other origins. While, in theory, algorithms are celebrated for their ability to outperform humans in computational ability, in practice they may present biases in their own right, insofar as their prediction does not correctly identify the right causality direction and neglects to account for all correlated variables (Blake et al., 2015). Furthermore, algorithms optimize choices based on the past history, thereby incorporating consumers' behavioral biases (Tucker, 2018; O'Neil, 2017).

In this paper we provide a systematic review of the economic implications of the new technologies on markets and society. By adopting both a macro and a micro perspective, we mean to capture a general picture of the recent but rapidly expanding economic literature on AI. The rest of the paper is organized as follows. Section 2 describes the effects of AI on growth, productivity and employment; Section 3 studies the effects on markets and competition. Finally, Section 4 concludes.

2 Socio-economic impacts of AI: growth, employment and public policies

In the last few years, remarkable progress has been made in machine learning, robotics and artificial intelligence applications (AI) and the technological change has already spread in many industries. However, scholars are still debating on how AI should be conceptualized, and its modeling often presents a strong connection with standard models of automation, seeing AI as a factor increasing the productivity of traditional inputs of production.

Moreover, from a policymaking point of view, some fundamental questions arise about the socio-economic returns of the new technologies. What is the expected impact of AI and ML on growth and productivity? What is their effect on the labor market? Finally, what is the role of public policies in avoiding mass job displacement and spiraling inequality out of control?

2.1 From modeling automation to modeling AI

Most of the literature on the effects of AI on economic growth focuses on how the new technologies could substitute labor in the production process. These models (see e.g. Kotlikoff and Sachs, 2012; Graetz and Michaels, 2015; Nordhaus, 2015) focus specifically on the potential of AI of increasing automation, and thus formalize it as a capital-augmenting factor (i.e., it increases output for each unit of capital used for the production). The assumption of a more productive capital limits the scope for labour and wage reductions, owing to the complementarity between capital and labour: anything that increases the productivity of

capital, increases also the marginal product of labour, and hence the labour demand. Under this lens, AI could serve as a new productivity engine for the global economy and it leads to higher economic growth by the channel of labor substitution. Graetz and Michaels (2015) estimate that robotics added a 0.4 percentage points of annual GDP growth between 1993 and 2007 on average for 17 developed countries, accounting for about one-tenth of GDP growth during the period considered.

The potential of new technologies for accelerating economic growth is studied by Nordhaus (2015), who explores the conditions under which AI would lead to an “economic singularity” and examines the empirical evidence on the elasticity of substitution on both the demand and supply sides of the economy. He argues that the rapid growth in computation capabilities and artificial intelligence will cross some boundary or “Singularity”, after which economic growth will accelerate sharply. The rapid economic growth can arise both from the demand and from the supply side, as a result of relatively elastic demand and production functions.

The fact that, as growth accelerates with superintelligent capital, the rate of return on capital and real interest rates fall to zero, was an outcome already envisioned by J.M. Keynes who foreboded the ‘euthanasia of the rentier’.¹ Keynes’s scenario described a growth path in which the elasticity of substitution between labor and capital is less than one; accumulation in the inelastic case therefore drives not only the rate of return to zero but also the share of capital to zero. However, the accelerationist case leads to the opposite outcome, where the share of capital goes to unity. In this outcome, we thus would see the euthanasia of the laboring classes, in the sense that all of income eventually goes to the owners of capital.

Should we be concerned about such an extreme vision? It might just be the case, but only if labour remains as it currently is. However, the impact of AI on the production process

¹In a chapter from *The General Theory* (Keynes, 1935 pp. 375-376), Keynes writes: “[There would be an] increase the stock of capital up to a point where its [marginal product] had fallen to a very low figure. . . . Now, [this] would mean the euthanasia of the rentier, and, consequently, the euthanasia of the cumulative oppressive power of the capitalist to exploit the scarcity-value of capital. Interest today rewards no genuine sacrifice, any more than does the rent of land. I see, therefore, the rentier aspect of capitalism as a transitional phase which will disappear when it has done its work. And with the disappearance of its rentier aspect much else in it besides will suffer a sea-change. It will be, moreover, a great advantage of the order of events which I am advocating, that the euthanasia of the rentier, of the functionless investor, will be nothing sudden, merely a gradual but prolonged continuance of what we have seen recently . . . and will need no devolution.”

is expected to have a larger scope than that of automation. A first step in this direction is made by Acemoglu and Restrepo (2018b), who adopt a different perspective and develop a model in which AI substitutes for workers in existing tasks but also creates new tasks for workers to do. They note that a distinctive feature of new technologies is the expansion of the set of tasks produced by machines, which creates the scope for the reduction of the demand of labour, unless the productivity gains from automation are sufficiently large. Building on Zeira's (1998) task-based approach, they assume that capital and labor are perfect substitutes in technologically automated tasks, but capital is cheaper. The role of AI is to expand the set of tasks where machines substitute labour, making the production less labour-intensive and reducing the labour share. Thanks to the perfect substitutability between capital and labour, the equilibrium wage decreases, despite the higher productivity. This result is more in line with the available empirical evidence about robots having a negative impact on local employment and wages (Acemoglu and Restrepo, 2017).

The task-based approach is also used by Aghion et al. (2018) to study the implications of AI for economic growth. They put AI in an automation framework, and argue that AI could be seen as the next progressive step of a two-centuries long automation process. They show that increased automation can have two effects on growth. First, the capital share of GDP increases, as automation is applied to a larger number of tasks. Second, the price of automated goods declines as a result of capital accumulation. If demand is relatively inelastic, the price effect prevails and the expenditure share of automated goods declines as well, thus explaining the periods of slow economic growth.

While these papers emphasize the potential substitution of intelligent machines for workers, Bessen (2018) argues that AI could also increase the productivity of labour, and he thus formalizes it as a labour-augmenting factor. He finds that new technologies should have a positive effect on employment if they improve productivity in markets characterized by substantial amount of unmet demand, such as non-manufacturing industries (Bessen, 2018).

In a similar vein, Agrawal, McHale and Oettl (2018) focus on the role of AI in supporting human researchers by improving the mechanisms of discovery in science. They incorporate artificial intelligence into an innovation-based growth model to show how AI can speed up growth by improving prediction accuracy and discovery rates.

The issue about the complementarity versus substitutability relationship between AI and labour is directly addressed by Agrawal, Gans and Goldfarb (2018). They consider a risky environment where a decision maker can choose between a risky or a safe action. AI reduces the cost of predictions and the decision maker can exercise human judgment, i.e. the ability to recognize hidden attributes of the venture. They show that a decision maker takes riskier actions either because he discovers hidden opportunities, or because the quality of predictions improves: hence, human judgment over hidden opportunities is a substitute of better predictions. Conversely, when prediction is precise, but the decision maker discovers some hidden cost, he reverts his decision to the safe action (i.e., human judgement on hidden costs is a complement of better predictions).

Departing from previous work, DeCanio (2016) considers AI is an entirely new factor of production. He breaks the link between AI and automation, by using a production function which explicitly accounts for labour, robots and ordinary capital, in order to study the substitution between human and robotic labour. He shows that an increase in robotic labor can have either a positive or a negative effect on wages, depending on the elasticity of substitution. The diffusion of robots will have a depressing effect on human wages only if the elasticity of substitution between human and robotic labor is relatively high.

2.2 Employment

Most of the concerns about the introduction of new technologies are about the adverse effects of automation on the labor market, such as which and how many jobs are going to be depleted. Acemoglu and Restrepo (2017) find that robot adoption has a significant and negative effect on employment in the US automotive sector. By making human work redundant and less competitive than capital, technologies result in a shift in the capital/labor mix and a change in the composition of labor demand.

Jobs involving repetitive, routine or optimization tasks are the ones most at risk of being replaced by intelligent machines. Conversely, jobs with greater creative or strategic content or that require social intelligence are less susceptible to computerization, although AI could assist people even in creative jobs or in those where empathy and human feelings play a

central role (Boden, 1998).

Fortunately, the effects of AI and ML on the labour market are not limited to their impact in terms of jobs destructions. On the positive side, the first positive consequence is that, as machine intelligence improves, the value of human judgment skills will increase. As human judgment is a complement to the machines' abilities, the demand and the value of judgment-related skills will rise (Agrawal, Gans and Goldfarb, 2016).

Moreover, AI will not only increase the value of jobs with a high content of human-related skills, it will also lead to the creation of new AI-driven business and technology jobs, as documented in a number of recent studies (Autor, 2015; Autor and Salomons, 2018; Brynjolfsson and McAfee, 2014). In this spirit, an extensive framework proposed by Acemoglu and Restrepo (2018c) considers both the automation of tasks that were previously executed using labor and the introduction of new tasks in which labor has a comparative advantage over capital. The main finding is that if the comparative advantage of labor over capital is sustainable and the number of the newly created tasks is sufficiently high, the demand of labor can remain stable (or even grow) over time, despite the process of automation. Dauth, Findeisen, Südekum and Wößner (2017) find that in Germany each additional industrial robot leads to the loss of two manufacturing jobs, but this loss is offset –and possibly over-compensated– by the creation of new jobs in the service industry.

This result, however, means that the demand for labor is constituted by skilled workers since the low-skill tasks can be easily automated and be performed by capital. The gap between low and high-skilled workers may rise, as unskilled labor is more substitutable to AI than the skilled one, leading to severe redistribution concerns (Tirole, 2017), as typically happens whenever new technologies are adopted (Ackerman et al., 2015).

This static framework does not consider the endogenous effects of automation on the price of capital. In Acemoglu and Restrepo (2018a), the long-run rental rate of capital relative to wages decreases due to capital accumulation. The authors show that, if the pressure is sufficiently strong, it could lead to automate all tasks.

The problem of income inequality is also center stage in the analysis of Hémous and Olsen (2016), who emphasize the implications for wage inequality between high-skilled and low-skilled workers. In their model, the introduction of machines substitutes low-skill labor

and complements high-skill labor, and it is combined with the introduction of new products, which increases demand for both types of labor. Investment in automation, made to meet the rising demand, depresses the future growth rate of low-skill wages, and reduces the total labor share. In the long run, low-skill wages grow but at a lower rate than high-skill wages.

2.3 Skills and organization

The effects of AI on labour market might be more complicated when the firms' internal organization is accounted for. Recent work by Aghion, Bergeaud, Blundell and Griffith (2017) suggests that while the prediction of a premium to skills may hold at the macroeconomic level, it perhaps misses important aspects of firms' internal organization. By using matched employer-employee data from the UK, they analyze the relationship between innovativeness and average wage income across firms, and find that lower-skilled workers benefit more from working in more R&D intensive firms (relative to working in a firm which does no R&D) than higher-skilled workers. This result could be explained by the fact that AI may increase the complementarity between low-skilled and high-skilled workers, which increases the bargaining power of low-skilled workers. In fact, the more innovative the firm, the more important it is to have high-ability low-occupation employees so as to make sure that the high-occupation employees within the firm concentrate on the most difficult tasks (Garicano and Rossi-Hansberg, 2006), hence the need to select out those low-occupation employees which are not trustworthy.

Within jobs, the introduction of advanced automation technologies might entail substantial reorganization of tasks. In this sense, Brynjolfsson et al. (2018) study the suitability of occupations for machine learning, and find that, while most occupations have at least some tasks that are suitable for automatization, only few of them can be fully automated. More importantly, unleashing the ML potential will require significant redesign of the task content of jobs. Therefore, the policymakers' concern should be also directed towards the redesign of jobs and the reengineering of business processes.

A further implication in terms of the internal organization of the firm, is that the introduction of AI allows firms to eliminate middle-range monitoring tasks, moving toward

flatter and decentralized organizational structures (e.g. see Bloom, Garicano, Sadun and Van Reenen, 2014). Cockburn, Henderson, and Stern (2018) argue that artificial intelligence technologies will likely affect also the organization of the innovation process.

At the same time, as AI technologies accelerate the number of tasks performed by machines and robots, greater skills will be needed by the humans who perform the remaining tasks, for both the efficient operation of firms as well as for utilizing AI and other technologies in the best possible way. Indeed, Makridakis (2017) argues that hiring, motivating and successfully managing talented individuals will be pivotal for a successful business strategy in the AI era, and it is a task which is nearly impossible to program into an algorithm.

AI also should encourage self-employment by making it easier for individuals to build up reputation (Tirole, 2017), and also through the outsourcing of low-occupation tasks. However, Tirole (2017) also explains that would be hasty to advocate the end of large corporations by AI, for two reasons. First, firms are better equipped than single individuals to bear the risks and costs of large fixed investments. Second, vertical integration facilitates relation-specific investments in situations of contractual incompleteness, which will reasonably persist despite the diffusion of AI.

2.4 Public policies

A key dimension for evaluating the potential of AI for disruption in the labour market is the speed with which innovations will take place. If changes will happen quickly, the economy is going to face sustained period of time in which large segments of the population are not working (see Goolsbee, 2017, for a discussion of speed of adoption and Acemoglu and Restrepo, 2016, for a useful model). These rapid changes, and the potential disruption to the workforce, suggest it is important that there are policies in place to support workers and their retraining.

In principle, the effects of the technological progress in terms of jobs destructions do not necessarily imply an efficiency loss. In fact, innovation by definition shifts out the production possibilities frontier, thus enabling the redistribution that is necessary to ensure that innovation leads to a Pareto improvement: the gains that arise to some factor owners as a

result of technological progress are excess returns that could be taxed away without introducing distortions into the economy (Korinek and Stiglitz, 2017). To address the problem of unemployment caused by the substitution of workers by machines, there has been a rising call for Universal Basic Income (UBI) policies, which would grant a minimal level of income to people regardless of their employment status. A UBI exhibits three fundamental features. First, it is available universally, i.e. open to all with the exclusion of non-citizens. Second, it typically provides cash transfers, although at least part of it could be transferred through in-kind benefits such as food vouchers, housing benefits, travel coupons. Third, it is unconditional, so as to satisfy the inclusiveness criterion regardless of the income or employment status of the applicant.

Although a UBI policy could be an efficient way to alleviate poverty, Goolsbee (2018) points out the challenges connected to its implementation on a large scale. First, a basic income would likely cause a sizable drop in labor market participation by low wage earners. Second, a broad-based UBI will likely shift money away from the poorest class of the population, which is that most adversely affected by the advent of the new technologies, with regressive effects on the economy. In fact, UBI is likely to reduce the net transfers currently received by households with lower incomes, larger number of children, or with a disability status, thus raising equity concerns.

An alternative to UBI is employment subsidies, which would lack the universality feature and be conditional on work. If, on the one hand, these policies could increase participation in the labor force (Eissa and Liebman, 1996; Hotz, Mullin, and Scholz, 2006), on the other hand they exhibit higher administrative costs relative to UBI, because of the need to verify the eligibility conditions. Moreover, they could also increase the incentive for fraud, for example misreporting earnings or hours. However, as the subsidies would phase out at higher incomes, the overall cost would be considerably smaller than UBI.

The net benefits of technological advances might be more nuanced when compared with sufficiently large market imperfections, such that first best lump sum transfers between workers and innovators are impossibly costly to implement. In this case, changes in patent length and capital taxation can act as a second-best device to redistribute surplus (Korinek and Stiglitz, 2017). In fact, there is a growing consensus that one of the sources of the

growth of inequality is the growth of rents earned by innovators (see e.g. Korinek and Ng, 2017). Taxing and redistributing such rents has an important role in ensuring that advances in technology are Pareto improving.

3 AI, digital markets and competition

The exploitation of AI technologies and the increasing use of algorithms has been described as a game changer (Ezrachi and Stucke, 2016) and it is expected to have a massive impact on existing markets. By reducing certain types of costs and by shifting the cost structure associated to the economic activity, it will significantly impact the functioning of the market, entailing both positive and negative effects on the consumers' surplus. Moreover, the increasing size of network effects is rapidly shifting the mode of competition, moving to more platform-oriented types of competition (Goolsbee, 2018).

3.1 Market efficiencies

The widespread use of algorithms is undoubtedly associated to significant efficiencies, which benefit firms as well as consumers. Some of these benefits originate from the use of digital technologies, which entail a reduction of search costs, replication costs, transportation costs, tracking costs and verification costs (Goldfarb and Tucker, 2019). However, the widespread use of algorithms is also associated to some specificities, relative to standard digital technologies.

3.1.1 Supply side efficiencies

On the supply side, algorithms can promote static efficiencies by reducing the cost of production, by improving the quality of existing products and by optimizing resource utilization and commercial strategies instantaneously following trials and feedback. For example, algorithms are being employed by insurance companies to better assess the risk of customers, make automatic offers, and even process claims. The Economist (2017) reports that once a customer was able to receive the reimbursement three seconds after he filed his claim on the

app. In those three seconds the machine reviewed the claim, ran 18 anti-fraud algorithms, approved it, sent payment instructions to the bank and informed the policyholder. Supply side efficiencies are also due to the fast-growing use of dynamic pricing. Dynamic pricing allows for instantaneous adjustment and optimization of prices based on many factors—such as stock availability, capacity constraints, competitors’ prices or fluctuations of demand. This guarantees that the market is constantly in equilibrium, preventing unsatisfied demand and excess of supply. Still, dynamic pricing strategies make it challenging for non-algorithmic sellers to compete and for consumers to make decisions under constant price fluctuations, unless they also use algorithms to facilitate decision-making.

Algorithms can also promote dynamic efficiency by triggering a virtuous mechanism whereby companies are under constant pressure to innovate (Cockburn, Henderson and Stern, 2018; OECD, 2015a). Indeed, algorithms have been used to develop new offerings, thus promoting market entry (OECD, 2016e; OECD, 2016f and OECD, 2016g).

Advanced tracking technologies allow firms to increase the effectiveness of advertising by adopting targeted ads (Goldfarb and Tucker 2011b, Goldfarb 2014). Although better targeting can increase competition between advertisers (Athey, Calvano, and Gans, 2018), Levin and Milgrom (2010) show that too much targeting can lead to insufficient competition among advertisers for the user attention sold by a monopolist media firm. Measuring the effectiveness of targeted advertising presents its own challenges because of a selection bias: given that targeted consumers are likely already aware of the advertised product, simple correlational research might overestimate the effectiveness of advertising. Indeed, Blake, Nosko, and Tadelis (2015), using data from a large field experiment at eBay, question the effectiveness of paid search ads by showing that targeted consumers would often visit the website anyway. On the other hand, Simonov, Nosko, and Rao (2018) use data from Microsoft’s Bing search engine to show that advertising might be more effective in the case of less well-known brands.

3.1.2 Demand side efficiencies

Algorithms can support consumer decisions by making information better organized and accessible. AI-powered search engines provide information on dimensions of competition other than prices, such as quality and consumers' preferences, so as to significantly reduce search and transaction costs and information asymmetries. Interestingly, however, lower search costs do not automatically translate into a consumer's benefit. As noted by Goldfarb and Tucker (2019), search costs are endogenous, and firms can manipulate the search process in order to escape the competitive pressure. Moreover, it should also be remarked that there are limits to the ability of consumers to process large amount of information, as consumers may suffer from information overload.

Algorithms could help overcome the problem of information overload by taking charge of the processing of information. Indeed, they can shift the decision-making process by allowing consumers to outsource purchasing decisions to algorithms, thereby originating the concept of "algorithmic consumer" (Gal and Elkin-Koren, 2017). In this way, algorithms help consumers to overcome behavioral biases and cognitive limits, make more rational choices and empowers them against manipulative marketing techniques.

Search engines can provide consumers' value not only by increasing the quantity of information, but also its quality, by delivering more relevant search results. However, the empirical evidence on the performance of search algorithms is so far scant and mixed. Chiou and Tucker (2017) find that the length of time that search engines retain their server logs has no significant impact on the accuracy of search results. Conversely, Bajari et al (2018), using data from Amazon, show that the predictions on sales of retail products improve by data that cover longer time periods (though at a diminishing rate), but not by data on larger number of products.

In addition, low-cost tracking technologies allow for personalized shopping recommendations. In principle, this should increase welfare by suggesting purchases that better match people's preferences (Brynjolfsson, Hu, and Smith, 2003). However, Fleder and Hosanagar (2009) demonstrate that recommendation engines that emphasize "people who bought this also bought" also reinforce the superstar effect (Rosen 1981), so that variety might decrease

endogenously.

Although popularity information is typically self-reinforcing (see, for example, Salganik et al. 2006, Cai et al. 2009, Zhang 2010, Chen et al. 2011), Tucker and Zhang (2011) show that it might positively affect consumers' surplus in the case of niche products. Using field experiment data from a website that lists wedding service vendors, they find that narrow-appeal vendors receive more visits than equally popular broad-appeal vendors, as popularity in a niche product is considered as a signal of quality.

Algorithms can also provide value to consumers indirectly. By making information more accessible, they increase the competitive pressure on firms, which are induced to compete on quality as well as prices. Moreover, algorithms can be usefully employed by Antitrust authorities as a detection tool to identify instances of coordination between suppliers and collusive pricing (OECD, 2017). Data-driven approaches have been proposed to detect bidding anomalies and suspicious bidding patterns across large data sets, particularly through the use of screening methods (OECD, 2017; Akhgar et al., 2016). For example, Akhgar et al. (2016) argue that machine learning algorithms could be applied to identify hidden relationships as an indicator of collusion in public tenders.

3.2 Market failures

The adoption of AI technologies on a large scale is potentially associated to negative effects too, which could result in a reduction of the consumers' surplus. As these advanced technologies typically entail relatively large fixed costs, there is greater scope for significant economies of scale. In many cases, these scale effects are also combined with important network externalities and large switching costs on the demand side of the industries (Varian, 2018). Because of network effects, it might be difficult for entrants to induce consumers to switch to their platform, even if data portability were possible. Data portability alone is insufficient to increase competition. Overall, all of these factors lead to situations of market power and a winner-take-all market structure (Cremer et al., 2019). Indeed, Bloom, Garicano, Sadun and Van Reenen (2014) argue that the introduction of AI may speed up the process by which each sector becomes concentrated over time. In what follows we discuss the

most immediate implications in terms of consumers' welfare and competition, which might call for the attention of Antitrust authorities.

3.2.1 Price discrimination

Firms' pricing decisions are increasingly delegated to AI-based algorithms (Chen et al., 2016). Algorithmic prices may be conditioned on a large number of variables, such as the timing of the purchase, the firm's residual capacity, but also on the consumer's entire past purchasing history. The larger availability of data and more sophisticated estimation methods help a finer targeting and segmentation of the market (Milgrom and Tadelis, 2018), dramatically enlarging the scope for price discrimination. First-degree price discrimination, so far only a theoretical possibility, could become a reality. Moreover, in an algorithm-driven environment, discrimination can be subtler than the classical price discrimination, and take the form of behavioral discrimination (Ezrachi and Stucke, 2016). Firms can harvest our personal data to identify which emotion (or bias) will prompt us to buy a product, or our reservation price. Advertising and marketing activities can be tailored to target us at critical moments with the right price and emotional pitch.

Despite the intense scrutiny of regulators and competition authorities to uncover such practices, few instances of first degree price discrimination have actually been observed in practice. The only empirical test of scalable price targeting is provided by Dubé and Misra (2017), who study its welfare implications by using a machine learning algorithm with a high-dimensional vector of customer features. In their study, they find that the firm's profits increase by over 10% under targeted pricing relative to the optimal uniform pricing, while overall customer surplus declines by less than 1%, although nearly 70% of customers are charged less than the uniform price. Shiller (2014) uses a Ordered-Choice Model Averaging Method to predict the subscription rates to Netflix. He shows that personalized prices based on the data about the web-browsing behavior of consumers –in addition to demographic variables- can significantly increase profits, while some consumers can pay as much as twice the price of others for the same product.

Gautier et al. (2019) argue that the scant evidence on AI-enabled personalized prices

can be attributed to technical barriers, as well as to several market constraints. First, price discrimination might not survive competition, especially when the competing firms share the same information about consumers (Belleflamme et al., 2017). Second, reputational concern of firms may limit the use of price discrimination, as consumers resent it as an exploitative practice. Third, consumers tend to react strategically to price discrimination by limiting the amount of information they reveal (Townley et al., 2017).

3.2.2 Other exclusionary conduct

Data are a crucial input for AI-based technologies. Therefore, ownership of large datasets may create barriers to entry and critically influence the efficient functioning of the competitive environment.

If new entrants are an important source of potential innovation, exclusionary conduct by incumbents can slow the pace of innovation (Chevalier, 2018). However, forcing the incumbent to share data could diminish the incumbent incentive to invest in data creation. These considerations clearly lead to antitrust enforcement concerns. Indeed, recent empirical work (e.g. see Aghion, Howitt and Prantl, 2015) points at patent protection and competition policy being complementary in inducing innovation and productivity growth.

Algorithms can facilitate exclusionary and exploitative anti-competitive practices also by the selection of search results (Patterson, 2013), whereas information is selected for business interests instead of its relevance or quality; Bar-Ilan (2007) finds that search engines are biased in order to rank certain websites higher. However, Ratliff and Rubinfeld (2014) warn that a consistent competitive assessment should account for the two-sided nature of the markets of search engines.

Manipulation can extend beyond search results. Online feedback is probably the most important reputation mechanism used for building trust in electronic markets. Dellarocas (2000) provide evidence of the creation of fake accounts in order to manipulate feedback scores and influence ratings. Thus, the digitalization of the word-of-mouth presents entirely new challenges from a policymaking point of view, including the definition of mechanisms for eliciting sufficient, honest and unbiased feedback, or for the creation of online identities

(Dellarocas, 2003).

3.2.3 Collusion

Algorithmic pricing may facilitate collusion via two main channels. First, algorithmic pricing allows to react to rivals' actions much more quickly than human beings (Ezrachi and Stucke, 2016; Mehra, 2016). Because of the frequent interactions, defection from a collusive agreement is punished more promptly and the gains from defection are reaped for a shorter time.

Second, last-generation algorithmic pricing is based on ML techniques. The software actively learns the optimal strategy purely by trial and error, by intentionally experimenting sub-optimal prices. This kind of pricing algorithms are highly flexible because they do not require the specification of the economic model as an input, and thus turn out to be particularly suitable in complex environments. Quite importantly, pricing algorithms might learn autonomously to set supra-competitive prices. Klein (2018) shows that simple algorithmic agents could learn to collude in a sequential move game. Moreover, in a recent experiment (Calvano et al., 2019), simple pricing algorithms are let to interact in a repeated pricing game with simultaneous moves and full price flexibility. The study finds that AI pricing agents systematically learn to play sophisticated collusive strategies, meting out punishments that are proportional to the extent of the deviations and are finite in duration, with a gradual return to the pre-deviation prices. Differently from collusion between human subjects, or to earlier experiments with Q-learning algorithms (Waltman and Kaymak, 2008), the collusive strategies played by AI agents are robust to perturbations of cost or demand, number of players, asymmetries and forms of uncertainty.

From an Antitrust standpoint, a critical problem is that the pricing algorithms leave no trace of concerted action – they learn to collude purely by trial and error, without communicating with one another, and without being specifically designed or instructed to collude. This poses a real challenge for competition policy, for two reasons. First, the current legal standard for collusion in most countries (including Europe and the US) has been designed for human agents, and thus requires some explicit intent and communication among firms to

restrain competition. Therefore, it fails in the case of tacit forms of collusion. Second, when pricing decisions are made by a machine using an algorithm rather than by human beings, establishing liability might be non-trivial and requires a revision of the current regulatory practices (OECD, 2017). Could liability be charged on the person who designed the algorithm, on the individual who used it or on the person (or entity) who benefited from the decision made by the algorithm, even if consumers' harm was not consciously done? The answer to such a question is not clear-cut at the moment. Gautier et al. (2019) observe that such a scenario might not ever materialize, as there are technical and market barriers that hinder the emergence of algorithmic tacit collusion outside the realm of lab experiments.

3.2.4 Privacy

The relative power of consumers over sellers is crucially affected by the regulation of privacy issues. Restrictions on consumer privacy and the ways that companies can use customer information can *de facto* be seen as an argument over property rights, in the sense of establishing who owns the consumers' data and what level of consent it requires to use it. Indeed, a central issue in terms of privacy is the extent of control that a consumer has not only on his personal information, but also on the information that can be inferred by algorithms by identifying patterns in his behavior.

Traditional models of information economics tend to think of consumer privacy as an information asymmetry and suggests that both buyers and sellers have an incentive to hide or reveal private information, crucially affecting market efficiency. In this context, less privacy is not necessarily bad for economic efficiency, as it could favor a better match between product and consumer type and it could help buyers to assess product quality, thus encouraging high quality production.

However, new technological advances have enabled a steep decline in the cost of collecting, storing, processing and using data in massive quantities, and thus extend information asymmetry far beyond a single transaction. Big data can be used as an input in order to understand, predict and influence consumer behavior long after the transaction has taken place. More importantly, there might be incomplete information about future data use, on

both sides of the transaction. As a consequence, the seller may be reluctant to restrict data use to a particular period, purpose or processing method.

Jin (2018) notes that AI exacerbates three problems related to consumers' privacy. First, sellers might have more information about future data use than buyers; as a consequence, sophisticated consumers hesitate to give away their personal data and they must trade-off between immediate gains from the transaction and potential loss from future data use. Second, sellers need not fully internalize potential harms to consumers because it is difficult to trace harm back to the origin of data misuse. Third, sellers have a higher incentive to renege on their consumer-friendly data policy, as it is difficult to detect and penalize it *ex post*.

Tucker (2017) argues that in terms of consumers' privacy, AI increases the potential of data persistence, data repurposing and data spillovers. In fact, once created, personal information may potentially persist longer than the human who created it, given the low costs of storing such data. Moreover, at the moment in which the data are created, there is uncertainty about how such data could be used in the future. There are also potential spillovers for others who did not provide the information, but are somehow affected by it.

3.2.5 Algorithmic biases

Automated computer systems are increasingly used to organize and select relevant information, such as the ordering of search results, the news that online users read, the multimedia content they access or the suggestions on future purchases. Such a role is not necessarily undesirable, especially because machines are considerably more efficient and objective than human beings in selecting relevant and quality information, potentially leading to better matching and reduced search costs. Nevertheless, computer selection biases might be present, leading to a whole new range of policy concerns. Algorithmic biases may be of different types (Saurwein et al., 2015), and occur for two main reasons: first, they make prediction based on data which is endogenously generated; second, they incorporate the behavioral biases of human beings.

Automated data-decision processes typically consider large amount of data, including

personal and demographic information. This may result in discriminatory –and potentially unintended- outcomes.

Recently, a heated debate arose regarding the use of algorithms for predicting recidivism in courtrooms. Angwin et al. (2016), analyzing the efficacy of the predictions on more than 7000 individuals arrested in Florida between 2013 and 2014, find that the software used was twice as likely to mistakenly flag black defendants as being at a higher risk of recidivism, and twice as likely to incorrectly flag white defendants as low risk. Although the data used by the algorithm do not include an individual’s race, other aspects of the data may be correlated to race that can lead to racial disparities in the predictions, thus opening a debate about the fairness criterion that should be used (Chouldechova, 2017).

Discrimination might also arise from crowding-out effects. Lambrecht and Tucker (2017) show that an ad for jobs in the Science, Technology, Engineering and Math fields is less likely to be shown to women, despite the fact that the ad is gender-neutral, and women are more likely to click on it - conditional on being shown the ad - than men. Moreover, the effect persists across 190 countries, so it does not depend on cultural factors. Interestingly, it appears that the algorithm is reacting to spillovers across advertisers. In fact, profit-maximizing advertisers pay more to show ads to females than males, especially in younger demographics, as the former often deliver a higher return on investment.

Algorithmic flaws might also originate from correlations in behavior. For example, Kosinski et al. (2013) report that someone liking (or disliking) ‘Curly Fries’ on Facebook is predictive of intelligence and it therefore could be used as a screening device by algorithms whose goal is to identify desirable employees or students.

Miller and Tucker (2018) find that an advertising algorithm tends to overpredict the presence of African Americans in states where there is a historical record of discrimination against African Americans. In fact, African Americans are more likely to have lower incomes in states which have exhibited historic patterns of discrimination (Sokoloff and Engerman, 2000; Bertocchi and Dimico, 2014). In turn, low-income people are more likely to use social media to express interest in celebrities movies, TV shows and music, as opposed to news and politics, which allows the algorithm to infer their ethnicity.

All these cases highlight the potential for historical persistence in algorithmic behavior,

which occurs because they make predictions on the basis of endogenously generated data (Tucker, 2019). Policymakers' awareness of these dynamics is necessary not to reinforce old, familiar biases and stereotypes. Mitchell and Brynjolfsson (2017) also note that algorithmic skews could be mitigated by integrating data from different sources.

In addition, algorithms can deploy information filters that reduce the variety and bias information according to the preferences of online users, leading to echo chambers (Sunstein, 2009) and filter bubbles (Pariser, 2011). For example, search engines provide readers with news that match their own beliefs and preferences. Product recommendations are biased towards similar content to previous purchases. Personalised content and services could limit the diversity of media content people are exposed to and thus have an adverse effect on the democratic discourse (e.g., Pariser, 2011; Sunstein, 2002; Viķe-Freiberga et al., 2013). However, they are worrisome also for other reasons. First, information filters are opaque and their criteria are invisible, hence it is difficult to form a belief about the extent to which the information received is biased. Second, with implicit personalization, people do not choose the filters and they might not even be aware of their existence, thus affecting how they respond to personalized messages (Viķe-Freiberga et al., 2013). Third, by limiting the exposure to diverse information, they constitute a centrifugal force of attitudinal reinforcement, making people drift towards more extreme viewpoints (Sunstein 2002, p. 9).

Algorithmic outcomes may turn out discriminatory also because the algorithm itself will learn to be biased on the basis of the behavioral data that feeds it (O'Neil, 2017). Documented alleged algorithmic bias spans charging more to Asians for test-taking prep software, to black names being more likely to produce 'criminal record' check ads (Sweeney, 2013), to women being less likely to seeing ads for an executive coaching service (Datta et al., 2015).

Probably, the largest scope of interaction between new technologies and people's behavioral responses is on the matter of privacy. Tucker (2018) notes that people could myopically reveal sensitive information that could harm them in the future, a problem aggravated by the persistence of data. John et al. (2011) find that individuals are less likely to provide personal information to professional-looking sites than unprofessional ones, or when they receive strong assurances that their data will be kept confidential. People might also suffer from an

illusion of control bias, that increases their willingness to share personal information when their perceived control over the release and access of information increases (Brandimarte et al., 2013).

People’s behavioral biases could be exploited to promote disclosure. Default settings, for example, can affect individual’s privacy behavior because they are not only convenient, but they are also often interpreted as implicit recommendations (McKenzie, et al., 2006). Gross and Acquisti (2005) find that only few individuals change the default privacy preferences, which are meant to maximize one’s profile’s visibility on social networks. Johnson et al. (2002) find evidence of the effects of opt-in or opt-out privacy policies on websites.

The increasing pervasiveness of computers calls for the understanding of how humans actually behave in interactions with intelligent machines, so as to properly guide the design of ICT systems. Johansen et al. (2016) propose that Behavioural Computer Science should integrate concepts from behavioral sciences into computer science models.

4 Conclusions

The implications of AI in terms of labour market outcomes has largely dominated the policy discussion in recent years, with economists highlighting both its job destruction and job creation potentials. AI certainly presents challenges in terms of wage inequality and unemployment, and policymakers will have a crucial role in determining the diffusion patterns and the impact of this technological revolution, as discussed by Agrawal et al. (2019). However, the economic effects of the recent technological advances go beyond their impact on the labour market. AI technologies can provide important and direct consumer benefits, through higher-quality and more accessible information. They will also have a massive impact on the functioning of existing markets, possibly redefining their notion and the ways with which firms interact between themselves and with consumers. This, however, might entail some threats for consumers’ welfare, increasing the risk of new elusive forms of collusion and firms’ exploitative practices. In this paper, we provide an overview of the many and multi-faceted effects of the recent technological advances, focusing on those issues with the most urgent policy implications.

Antitrust and privacy authorities alike will face unprecedented challenges to face the new complex and rapidly evolving environment. First, the access to data may act as an entry barrier for creating new competing networks and for investing in innovation by new market participants; this will also increase the incentive to undertake anticompetitive conduct in non-price dimensions, like data capture, extraction and exclusion. Second, the increased ability to track individuals enables novel forms of price discrimination. Third, quite importantly, the use of AI technologies is expected to widen instances in which known forms of anticompetitive conduct occurs, such as express and tacit collusion and discrimination (Petit, 2017). The use of advanced machine learning algorithms is likely to increase the opacity of the price setting process adopted by firms, thereby making it challenging for antitrust authorities to detect and punish anticompetitive conduct. Fourth, the use of massive quantities of data by AI technologies raises the risk of data manipulation, with important implications from a social and political point of view. For example, the control over search results can also be exploited for political interests. Epstein and Robertson (2015) show that biased search rankings can shift the voting preferences of undecided voters by 20%, with people not even being aware of the manipulation. While extremely relevant, the issue of data agglomeration and exploitation for political purposes is beyond the scope of this survey.

Although AI-based tools may provide a precious support to policymakers and improve policy accuracy, there are limits to the scope of their action. As Goolsbee (2018, p.8) puts it, “the technology may improve our ability to predict responses, but it does not help us balance interests or engage in politics”.

References

- [1] Acemoglu, D. and Restrepo, P. (2017). Robots and Jobs: Evidence from US Labor Markets. NBER Working Paper 23285.
- [2] Acemoglu, D. and Restrepo, P (2018a). The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment. *American Economic Review*, 108(6): 1488-1542.

- [3] Acemoglu, Daron, and Pascual Restrepo. 2018b. “Modeling Automation.” National Bureau of Economic Research Working Paper 24321.
- [4] Acemoglu, D. and Restrepo, P. (2018c). Artificial Intelligence, Automation and Work. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- [5] Acquisti, A., Taylor, C. and Wagman L. (2016). ”The Economics of Privacy.” *Journal of Economic Literature*, 54 (2): 442-92.
- [6] Akerman, A., Gaarder, I. and Mogstad, M. (2015) The Skill Complementarity of Broadband Internet, *The Quarterly Journal of Economics*, Volume 130(4): 1781-1824.
- [7] Aghion, P., Bergeaud, A., Blundell, R. and Griffith, R. (2017). The innovation premium to low skill jobs. Mimeo.
- [8] Aghion, P., Howitt, P. and Prantl, S. (2015). Patent rights, product market reforms, and innovation. *Journal of Economic Growth*, 20(3), 223–262.
- [9] Aghion, P., Jones, B. and Jones, C. (2018). Artificial Intelligence and Economic Growth. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- [10] Agrawal, A., Gans, J. and Goldfarb, A. (2016). The simple economics of machine intelligence. *Harvard Business Review* 17.
- [11] Agrawal, A., Gans, J. and Goldfarb, A. (2018). Human Judgement and AI pricing. NBER working paper 24284
- [12] Agrawal, A., Gans, J. and Goldfarb, A. (2019). Economic Policy for Artificial Intelligence, *Innovation Policy and the Economy* 19: 139-159.
- [13] Agrawal, A., McHale, J. and Oettl, A. (2018). Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

- [14] Akhgar, B., P. S. Bayerl and F. Sampson (2016), *Open Source Intelligence Investigation: From Strategy to Implementation*, Springer International Publishing.
- [15] Angwin J, Larson J, Mattu S, Kirchner L. (2016). Machine bias: There's software used across the country to predict future criminals and it's biased against blacks. Available online at www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing (last accessed June 4, 2019).
- [16] Athey, S., Calvano, E. and Gans, J.S. (2018). The Impact of Consumer Multi-homing on Advertising Markets and Media Competition. *Management Science* 64(4): 1574-90.
- [17] Autor, D. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation, *Journal of Economic Perspectives*, 29(3), 3–30.
- [18] Autor, D. and Salomons, A. (2018). Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share. NBER Working Paper 24871.
- [19] Bajari, P., V. Chernozhukov, A. Hortaçsu, J. Suzuki (2018). The Impact of Big Data on Firm Performance: An Empirical Investigation. National Bureau of Economic Research working paper No. 24334.
- [20] Bar-Ilan, J. (2007), Google Bombing from a Time Perspective, *Journal of Computer-Mediated Communication*, Vol. 12, No. 3, pp. 910-938.
- [21] Belleflamme, P., Lam, W. and Vergote, W. (2017). Price discrimination and dispersion under asymmetric profiling of customers, CORE discussion paper, Louvain-la-Neuve, Belgium.
- [22] Bessen, J. (2017). Automation and Jobs: When Technology Boosts Employment, Boston Univ. Law and Economics Research Paper No. 17-09.
- [23] Bessen, J. (2018). AI and Jobs: the Role of Demand. NBER Working Paper no. 24235
- [24] Blake, T., Moshary, S., Sweeney, K. and Tadelis, S. (2018). Price Salience and Product Choice. National Bureau of Economic Research Working Paper 25186.

- [25] Blake, T., Nosko, C. and Tadelis, S. (2015). Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment. *Econometrica* 83 (1): 155–74.
- [26] Bloom, N., Garicano, L., Sadun, R. and Van Reenen, J. (2014). The Distinct Effects of Information Technology and Communication Technology on Firm Organization. *Management Science*, 60(12): 2859-2885.
- [27] Boden M.A. (1998). Creativity and Artificial Intelligence. *Artificial Intelligence* 103: 347-356.
- [28] Bresnahan, T.F. and M. Trajtenberg (1995). General purpose technologies ‘Engines of growth’?, *Journal of Econometrics*, 65(1): 83-108.
- [29] Brandimarte, L., Acquisti, A., Loewenstein, G., (2013). Misplaced confidences privacy and the control paradox. *Soc. Psychol. Personal. Sci.* 4, 340–347
- [30] Brynjolfsson, E., Y. Hu, and M. D. Smith (2003). Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers. *Management Science* 49 (11): 1580–96.
- [31] Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York, NY: W.W. Norton, 2014.
- [32] Brynjolfsson, E., Mitchell, T. and Rock, D. (2018). What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 108: 43-47.
- [33] Calvano, E., Calzolari, G., Denicolò, V. and Pastorello, S. (2019), Artificial intelligence, algorithmic pricing and collusion, CEPR Discussion Paper 13405.
- [34] Chen, L., Mislove, A. and Wilson, C. (2016). An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace. In *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, pp. 1339–1349.

- [35] Chevalier, J. (2018). Antitrust and Artificial Intelligence: Discussion of Varian. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- [36] Chiou, L. and C. E. Tucker. (2017). “Search Engines and Data Retention: Implications for Privacy and Antitrust.” NBER Working Paper No. w23815.
- [37] Chouldechova, A. (2017). Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 5(2): 153-163
- [38] Cockburn, I., Henderson, R. and Stern, S. (2018). The impact of Artificial Intelligence on Innovation. NBER Working Paper no. 24449.
- [39] Cremer, J., de Montjoye, Y-A., Schweitzer, H. (2019). Competition policy for the digital era. Final report for the European Commission, Directorate-General for Competition.
- [40] Dauth, W., Findeisen, S., Südekum, J. and Wößner, N. (2017). German robots –The Impact of Industrial Robots on Workers. IAB discussion paper.
- [41] DeCanio, S.J. (2016). Robots and Humans – complements or substitutes? *Journal of Macroeconomics*, 49: 280-291.
- [42] Dellarocas, C. (2003) The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms. *Management Science* 49(10): 1407-1424.
- [43] Dellarocas, C. (2000). Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. *Proc. 2nd ACM Conf. Electronic Commerce*. Association for Computing Machinery, Minneapolis, MN, 150–157.
- [44] Dubé J.-P. and Misra, S. (2017). Scalable Price Targeting. NBER Working Paper no. 23775
- [45] Eissa, N. and J. B. Liebman. 1996. “Labor Supply Response to the Earned Income Tax Credit.” *Quarterly Journal of Economics* 111(2): 605-37.

- [46] Epstein, R. and Robertson R.E. (2015). The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences*, 112 (33) E4512-E4521.
- [47] European Commission (2017). CASE AT.39740 Google Search (Shopping). Antitrust Procedure Council Regulation (EC) 1/2003. Available at [http : //ec.europa.eu/competition/antitrust/cases/dec_docs/39740/39740_149963.pdf](http://ec.europa.eu/competition/antitrust/cases/dec_docs/39740/39740_149963.pdf), last accessed July 4th, 2019.
- [48] Ezrachi, A. and Stucke, M. E. (2016). Virtual Competition. *Journal of European Competition Law & Practice*, 7 (9), 585–586.
- [49] Fleder, D., and K. Hosanagar (2009). Blockbuster Culture’s Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science* 55 (5): 697–712.
- [50] Gal, M.S. and N. Elkin-Koren (2017), Algorithmic Consumers, *Harvard Journal of Law and Technology*, Vol. 30(2): 309-352.
- [51] Garicano, L. and Esteban Rossi-Hansberg, “Organization and Inequality in a Knowledge Economy,” *The Quarterly Journal of Economics*, 2006, 121 (4), 1383–1435.
- [52] Gautier, A., Ittoo, A. and Van Cleynenbreugel, P. (2019). AI Algorithms, Price Discrimination and Collusion: A Technological, Economic and Legal Perspective. Mimeo.
- [53] Goldfarb, A. (2014). What Is Different about Online Advertising? *Review of Industrial Organization*, 44(2): 115-29.
- [54] Goldfarb, A. and Tucker, C. (2011b). Online Advertising. *Advances in Computers* 81: 289-315.
- [55] Goldfarb, A. and Tucker, C. (2019). Digital Economics. *Journal of Economic Literature*, 57(1): 3-43.
- [56] Goolsbee A. (2018). Public policy in an AI economy. NBER Working Paper 24653

- [57] Graetz, Georg and Guy Michaels (2015). Robots at Work, CEP Discussion Paper No 1335.
- [58] Green Leigh, N. and Kraft, B.R. (2017). Emerging robotic regions in the United States: insights for regional economic evolution. *Regional Studies*, 1-13.
- [59] Gross, R., Acquisti, A. (2005). Information Revelation and Privacy in Online Social Networks (ACM Workshop–Privacy in the Electronic Society, New York), pp. 71–80.
- [60] Hemous, D. and Olsen, M. (2016). The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality. University of Zurich manuscript.
- [61] Hossain, T., and Morgan, J. (2006). ...Plus Shipping and Handling: Revenue (non) Equivalence in Field Experiments on eBay. *Advances in Economic Analysis and Policy* 5(2).
- [62] Hotz, V. Joseph, Charles H. Mullin, and John Karl Scholz. 2006. “Examining the Effect of the Earned Income Tax Credit on the Labor Market Participation of Families on Welfare.” National Bureau of Economic Research working paper No. 11968.
- [63] Jin, G.Z. (2018). Artificial Intelligence and consumer privacy. NBER Working Paper No. 24253.
- [64] Johansen, C. , Pedersen, T. and Jøsang, A. (2016). Towards Behavioural Computer Science. 10th IFIP International Conference on Trust Management (TM), Jul 2016, Darmstadt, Germany. pp.154- 163.
- [65] John, L.K., Acquisti, A., Loewenstein, G. (2011). Strangers on a plane: Context-dependent willingness to divulge sensitive information. *J. Consum. Res.* 37, 858–873.
- [66] Johnson, E.J., Bellman, S., and Lohse, G.L. (2002). Defaults, framing and privacy: Why opting in \neq opting out. *Mark. Lett.* 13, 5–15
- [67] Keynes, J. M., (1930). Economic possibilities for our grandchildren. <http://www.econ.yale.edu/smith/econ116a/keynes1.pdf> (last accessed June 18th, 2016).

- [68] Keynes, J. M. (1935). *General theory of employment, interest and money*. Macmillan, London.
- [69] Klein, T. (2018). *Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing*. Amsterdam Law School Research Paper No. 2018-15; Amsterdam Center for Law & Economics Working Paper No. 2018-05.
- [70] Korinek, A. and Stiglitz, J.E. (2017). *Artificial Intelligence and Its Implications for Income Distribution and Unemployment*. NBER Working Paper No. 24174.
- [71] Korinek, Anton and Ding Xuan Ng. 2017. "The Macroeconomics of Superstars." Working paper.
- [72] Kosinski, M., D. Stillwell, and T. Graepel (2013). *Private traits and attributes are predictable from digital records of human behavior*. *Proceedings of the National Academy of Sciences* 110(15), 5802–5805.
- [73] Kotlikoff, L. and Sachs, J.D. (2012). *Smart Machines and Long-Term Misery*, NBER Working Paper No. 18629.
- [74] Lambrecht, A. and C. Tucker (2017). *Algorithmic discrimination? apparent algorithmic bias in the serving of stem ads*, Mimeo, MIT.
- [75] Levin, J., and Milgrom, P. (2010). *Online Advertising: Heterogeneity and Conflation in Market Design*. *American Economic Review*, 100(2): 603-07.
- [76] Lynch, J.G., and Ariely, D. (2000). *Wine Online: Search Costs Affect Competition on Price, Quality, and Distribution*. *Marketing Science* 19 (1): 83-103.
- [77] Makridakis, S. (2017). *The forthcoming Artificial Intelligence (AI) revolution: Its impact on society and firms*. *Futures*, 90: 46-60.
- [78] McKenzie, C.R., Liersch, M.J., and Finkelstein, S.R. (2006). *Recommendations implicit in policy defaults*. *Psychol. Sci.* 17, 414–420.

- [79] Mehra, S. K. (2015), Antitrust and the Robo-Seller: Competition in the Time of Algorithms, Minnesota. Law Review, Vol. 100.
- [80] Milgrom, P.R. and Tadelis, S. (2018). How Artificial Intelligence and Machine Learning Can Impact Market Design. NBER Working Paper 24282.
- [81] Miller, A. and C. Tucker (2018). Historic patterns of racial oppression and algorithms. Mimeo, MIT.
- [82] Mitchell, T. and Brynjolfsson, E. (2017). Track how technology is transforming work. Nature, 544(7650): 290-292.
- [83] Nordhaus, W. (2015). Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth, NBER Working Paper No. 21547.
- [84] OECD (2015a), Data-Driven Innovation: Big Data for Growth and Well-Being, OECD Publishing, Paris, <http://dx.doi.org/10.1787/9789264229358-en>.
- [85] OECD (2016e), Competition and Innovation in Land Transport, [https://one.oecd.org/document/DAF/COMP/WP2\(2016\)6/en/pdf](https://one.oecd.org/document/DAF/COMP/WP2(2016)6/en/pdf).
- [86] OECD (2016f), Protecting and Promoting Competition in Response to ‘Disruptive’ Innovations in Legal Services, [https://one.oecd.org/document/DAF/COMP/WP2\(2016\)1/en/pdf](https://one.oecd.org/document/DAF/COMP/WP2(2016)1/en/pdf).
- [87] OECD (2016g), Refining Regulation to Enable Major Innovations in Financial Markets, [https://one.oecd.org/document/DAF/COMP/WP2\(2015\)9/en/pdf](https://one.oecd.org/document/DAF/COMP/WP2(2015)9/en/pdf).
- [88] OECD (2017), Algorithms and Collusion: Competition Policy in the Digital Age www.oecd.org/competition/algorithms-collusion-competition-policy-in-the-digital-age.htm
- [89] Pariser, E. (2011). The filter bubble: What the Internet is hiding from you. New York, NY: Penguin.

- [90] Patterson, M.R. (2013), Google and search-engine market power, Harvard Journal of Law & Technology Occasional Paper Series.
- [91] Petit, N. (2017). Antitrust and Artificial Intelligence: A Research Agenda. *Journal of European Competition Law & Practice*, vol. 8, No. 6
- [92] Ratliff, J.D. and Rubinfeld, D-L. (2014). Is there a market for organic search engine results and can their manipulation give rise to antitrust liability? *Journal of Competition Law & Economics*, 10(3): 517–541.
- [93] Rosen, S. (1981). The Economics of Superstars. *American Economic Review* 71 (5): 845-58.
- [94] Saurwein, F., N. Just and M. Latzer (2015), *Governance of Algorithms: Options and Limitations*, Vol. 17(6), 35-49.
- [95] Shiller B.R. (2014) First degree price discrimination using big data. Economics Department, Brandeis University, MA.
- [96] Simonov, A., Nosko, C. and Rao, J.M. (2018). Competition and Crowd-Out for Brand Keywords in Sponsored Search. *Marketing Science* 37 (2): 200–215.
- [97] Sunstein, C. (2009), *Republic.com 2.0*, Princeton, NJ: Princeton University Press. Tirole, J. (2017). *Economics for the Common Good*, Princeton University Press, 2017.
- [98] *The Economist* (2017). When Life Throws You Lemons: A New York Startup Shakes Up the Insurance Business, www.economist.com/finance-and-economics/2017/03/09/a-new-york-startup-shakes-up-the-insurance-business. Last accessed July 4th, 2019.
- [99] Townley, C., Morrison, E. and Yeung, K. (2017). Big data and personalized price discrimination in EU competition law. *Yearbook of European Law*, 36: 683-748
- [100] Tucker, C. (2018). Privacy, Algorithms and Artificial Intelligence. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

- [101] Tucker, C.E., and J. Zhang (2011). How Does Popularity Information Affect Choices? A Field Experiment. *Management Science* 57 (5): 828–42.
- [102] Varian, H. (2018). Artificial Intelligence, Economics, and Industrial Organization. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press
- [103] Viķe-Freiberga, V., Däubler-Gmelin, H., Hammersley, B., Pessoa Maduro, L.M.P. (2013). A free and pluralistic media to sustain European democracy.
- [104] Waltman, L., and Kaymak, U. (2008). Q-learning agents in a Cournot oligopoly model. *Journal of Economic Dynamics and Control*, 32(10): 3275-3293.
- [105] Zeira, J. (1998). Workers, Machines, and Economic Growth. *Quarterly Journal of Economics* 113 (4): 1091–117.

Author contacts:

Laura Abrardi

Politecnico di Torino, DIGEP

Email: laura.abrardi@polito.it

Carlo Cambini

Politecnico di Torino, DIGEP

Email: carlo.cambini@polito.it

Laura Rondi

Politecnico di Torino, DIGEP

Email: laura.rondi@polito.it