

Artificial intelligence, firms and consumer behavior: A survey

Original

Artificial intelligence, firms and consumer behavior: A survey / Abrardi, L.; Cambini, C.; Rondi, L.. - In: JOURNAL OF ECONOMIC SURVEYS. - ISSN 1467-6419. - (2021). [10.1111/joes.12455]

Availability:

This version is available at: 11583/2919852 since: 2021-09-01T08:55:26Z

Publisher:

John Wiley and Sons Inc

Published

DOI:10.1111/joes.12455

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

ARTICLE

Artificial intelligence, firms and consumer behavior: A survey

Laura Abrardi¹  | Carlo Cambini^{1,2}  | Laura Rondi¹ 

¹ Politecnico di Torino, Department of Management, corso Duca degli Abruzzi 24, Torino, Italy

² Florence School of Regulation and European University Institute

Correspondence

Laura Abrardi, Politecnico di Torino, Department of Management.
Email: laura.abrardi@polito.it

Funding information

Ministero dell'Istruzione, dell'Università e della Ricerca, Grant/Award Number: TESUN-83486178370409 finanziamento Dipartimenti di Eccellenza CAP. 1694 TIT. 232 ART. 6

Abstract

The current advances in Artificial Intelligence (AI) 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 impact of these technologies on the labor market and firms? Will algorithms reduce consumers' biases or will they rather originate new ones? How competition will be affected by AI-powered agents? This study is a first attempt to survey the growing literature on the multi-faceted economic effects of the recent technological advances in AI that involve machine learning applications. We first review research on the implications of AI on firms, focusing on its impact on labor market, productivity, skill composition and innovation. Then we examine how AI contributes to shaping consumer behavior and market competition. We conclude by discussing how public policies can deal with the radical changes that AI is already producing and is going to generate in the future for firms and consumers.

KEYWORDS

Artificial Intelligence, algorithms, machine learning

JEL CLASSIFICATION

D24, E24, L50, L86, O33

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Journal of Economic Surveys* published by John Wiley & Sons Ltd

1 | INTRODUCTION

Artificial Intelligence (AI) plays an increasingly important role in our economy. AI has the potential to become the engine of productivity and economic growth. It can increase the efficiency and quality of decision-making processes and spawn the creation of new products and services, markets and industries. However, AI may also have detrimental effects on the economy and society. For instance, it entails serious risks of job market polarization, rising inequality, structural unemployment and emergence of new undesirable industrial structures. Policymakers need to create the conditions necessary for nurturing the potential of AI while considering carefully how to address the risks it involves.

In this paper we provide a systematic review of the most recent economic literature on AI and data-enabled machine learning and their impact on firms, consumers and markets, focusing on those aspects of the new technologies that pose imminent policymaking challenges. Our goal is to provide a unified picture of the potential impact of AI on different microeconomic dimensions, pointing out how public policies can deal with the radical changes associated with AI technologies. By doing so, we aim at providing a new perspective that could help scholars to move ahead from the fragmented state of the recent, but rapidly expanding literature and to develop more cumulative knowledge, particularly in those directions that are more conducive to inform the policy debate.

The literature on the economic effects of AI is nascent but rapidly growing. One of the main challenges of navigating this literature is the lack of a common framework to analyze AI (Agrawal et al., 2019a) and different approaches have been proposed. A first view conceptualizes AI as a predictive technology based on Machine Learning (ML). Through ML, a branch of computational statistics, AI systems can produce new knowledge by finding complex structures and patterns in example data (Taddy, 2019). Computer systems can perform tasks such as understanding natural language, diagnosing diseases and even driving a car. Due to the versatility of the technology, it is difficult to pinpoint a precise definition of AI. The OECD offers the following definition, that is specific enough to fit existing technologies but at the same time allows for policy implementation: “An AI system is a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments”.¹ The ability to learn with varying levels of autonomy distinguishes ML-based AI systems from earlier digital technologies. However, the ML-based vision of AI also emphasizes some limitations of current AI systems. First, ML technologies can only predict a future that follows the same patterns of past data (Taddy, 2019). Second, while AI is likely a substitute for human prediction, it still requires human skills such as judgment –i.e., the ability to define utility or valuation functions (Agrawal et al., 2018). Finally, AI systems still need human expertise to organize ML applications within a structure that is business-specific, and thus requires human knowledge (Taddy, 2019). Because of these limitations, supporters of the ML-based concept of AI question the possibility that AI will completely substitute human intelligence.

Alongside the ML-based approach, which focuses on the ability of machines to learn and make predictions, other concepts of AI have been proposed. A more prudent approach to AI feels that humans will still be better at thinking outside the box for many years to come (see, e.g., Boden, 1998). According to this vision, AI will mostly be used for labor augmentation, providing humans with insights, advice and guidance to increase firm’s productivity. This concept of AI presents strong similarities with the concept of automation, being a factor that increases the productivity of traditional inputs of production.² In contrast to this approach, a more futuristic-looking vision

foresees a general AI capable of out-performing human intelligence under any aspect (Bostrom, 2014; Kaplan, 2016). Then, society will have to deal with what has been defined as an economic singularity (Nordhaus, 2020): an economy of radical abundance characterized by unbounded growth, in which no one will need to work anymore. In this case, 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 & Stiglitz, 2019).

Regardless of whether AI is viewed as a predictive technology, as automation or as general machine intelligence, there is substantial agreement that it will have a relevant impact on our economy. However, trying to analyze what exactly this impact will be (and the role of public policy) through multiple lenses may generate ambiguity unless some boundaries are defined.

In this survey, we thus adopt the ML-based definition of AI and focus on those studies where AI involves some measure of data-enabled learning by machines. In the literature that we explore, intelligent machines can produce new knowledge, and not simply perform existing tasks more efficiently and accurately. This stance excludes, therefore, the very interesting literature on automation and robotics, or research where AI is cast within a standard capital-labor productivity framework as a capital-augmenting factor (Graetz & Michaels, 2018; Kotlikoff & Sachs, 2012; Nordhaus, 2020). We instead focus on studies that regard AI as a completely new input of production, – for example, supporting economic growth (Acemoglu & Restrepo, 2018a; Aghion, Jones, & Jones, 2019), or new decision-making tools (Athey et al., 2018; Calvano et al., 2020), and we single out papers that explore the policy implications, while leaving the reader to the book by Agrawal et al. (2019a) for a general analysis of AI and economics, to the survey of Goldfarb and Tucker (2019) for a review of the literature on digital economics, and to the work of Lu and Zhou (2021) for a review on the economics of AI from a macroeconomic perspective.

The repercussions of AI in the labor market, on consumers' behavior and on competition appear to have dominated the policy discussion so far. 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 et al., 2019b). Solving them will require concerted actions of redistribution of wealth and will possibly entail devising ways to train people to work on other tasks (Korinek & Stiglitz, 2019). However, what exactly this impact will be is an open issue, given that few empirical studies possess the granularity of data to differentiate AI and machine learning from industrial robots and automation.

AI is 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. Indeed, AI systems are challenging the current framework of the market mechanisms and of the consumers' decision-making processes. The enhanced role of data, the very strong economies of scale and scope and extreme network effects give rise to a strong incumbency advantage, and lead to highly concentrated markets with few dominant players (Cremer et al., 2019). Moreover, ML algorithms may generate distortions *per se* (Blake et al., 2015), or exacerbate consumers' behavioral biases (Tucker, 2019), that can be exploited by dominant digital players to further affect the market efficiency. Understanding such breakthroughs and their effects in terms of competition policy, privacy and the efficient allocation of resources (including data) is paramount to exploit the benefits and tackle the threats of the new technologies.³

The rest of the paper is organized as follows. Section 2 describes the effects of AI on firms, specifically in terms of productivity, labor, employment, and skill composition. Section 3 discusses the implications of algorithmic distortions on consumers and their behavior. Section 4 studies the effects on markets and competition. Finally, Section 5 offers some conclusions and ideas for future

research. To grasp a better picture of this analysis, we summarize the main contributions on AI in tables reporting the type of data and methodology used in each paper, as well as their main results.

2 | AI AND FIRMS

ML-driven AI is often applied as a general-purpose technology (Bresnahan & Trajtenberg, 1995), as it can be employed transversally across sectors. Indeed, technological change has already spread in many industries and this has raised the policymakers' concern about the impact of new technologies on the labor market (Agrawal, Gans, & Goldfarb, 2019b). In this section, we explore specifically their effects on employment, skill composition and firm organization of the innovation process.

2.1 | Labor and skills

Most of the concerns about the introduction of new technologies are related to their adverse effects on the labor market, such as which and how many jobs are going to be depleted. 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).⁴

Among this literature, there is a wide consensus that, because of AI, the demand for labor will increasingly be directed to skilled workers, since the low-skill, routine tasks can be easily performed by machines, leading to severe redistribution concerns (Tirole, 2017), as typically happens whenever new technologies are adopted (Akerman et al., 2015). As a consequence, public policies directed at the redistribution of wealth will likely play a central role in the near future, although subsidy-based policies present their own challenges. Policies like the Universal Basic Income, granting a minimal level of income to people regardless of their employment status, are costly to implement on a large scale, they might reduce the labor market participation by low wage earners, and they could also have regressive effects, as they are likely to shift money away from the poorest class of the population (Goolsbee, 2018). On the contrary, employment subsidies would increase the participation in the labor force (Eissa & Liebman, 1996; Hotz et al., 2006), but they entail significant administrative costs, because of the need to verify the eligibility conditions. Moreover, with sufficiently large market imperfections, fixed transfers to redistribute surplus from innovators to workers are impossibly costly to implement and other policies, such as changes in patent length and capital taxation, should be considered as a second-best device to redistribute surplus (Korinek & Stiglitz, 2019).

The concern for low-skill occupations is partly derived by past experience with robotics and software. However, recent research (Acemoglu et al., 2020; Webb, 2020) suggests that AI might have a different pattern to robots and software, as it greatly affects also high-skill, high-tech jobs. High-skilled occupations that require a college degree and accumulated experience are more likely to involve tasks like detecting patterns, making judgments, and optimization, that can be successfully automated by AI.

The negative effects of AI on the labor market have understandably dominated the policy debate. However, when the role of AI as a prediction technology is accounted for, its effects on the labor market are more nuanced (Agrawal et al., 2019a) and not limited to their impact

in terms of jobs' destruction. On the one hand, a substitution effect may still arise, as AI may directly substitute capital for labor in prediction tasks, and even in some decision tasks (specifically, when automating prediction increases the relative returns to capital versus labor), raising an issue of organizational design related to the optimal allocation of decision authority to the human rather than to the machine (Athey et al., 2020). On the other hand, AI might enhance labor when automating the prediction tasks, thus increasing labor productivity. Furthermore, AI may create new decision tasks, as long as better predictions sufficiently reduce uncertainty and enable new decisions that were not feasible before.

Highlighting the potential complementarity of AI with labor, Bessen (2018) argues that AI could increase labor productivity. Along these lines, Agrawal et al. (2016) suggest that human activities can be described by five high-level components: data, prediction, judgment, action, and outcomes. Judgment is making decisions based on prediction outputs by weighting options and payoffs. As machine intelligence improves, many tasks can be reframed as prediction problems, and the value of human prediction decreases, substituted by machine prediction. However, this process will increase the value and the demand of human judgment skills, which is a complement to the machines' abilities.

The issue about the complementarity versus substitutability relationship between AI and labor is directly tackled by Agrawal et al. (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, that is, 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 judgment on hidden costs is a complement of better predictions).

This literature highlights that, because of its complementarity with some human tasks, AI might increase the value of jobs with a high content of human-related skills. A further positive effect of AI is the creation of new AI-driven business and technology jobs, as documented in a number of recent studies (Autor, 2015, 2018; Brynjolfsson & McAfee, 2014). In this spirit, Acemoglu and Restrepo (2019) consider 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 for labor can remain stable (or even grow) over time.

Despite the progress of the theoretical literature on AI, few studies examine its impact from an empirical point of view. Recent work by Webb (2020) comparing job descriptions (from the O*NET database, the US government list of work activities and occupations) to patent descriptions identifies the most exposed occupations to AI⁵. He finds that most exposed occupations involve prediction tasks, optimization, and analytical work. The least exposed occupations instead involve interpersonal skills (such as teachers and managers), reasoning about situations that have never been seen before (e.g., researchers), or manual work that occurs in a non-factory environment (baristas, massage therapists). Descriptive evidence on the occupational impact of AI shows that workers in the 90th wage percentile are most exposed to AI. Webb (2020) also estimates through simulations that AI might reduce by 4% the 90:10 wage inequality, namely the ratio of the 90th to the 10th percentile of wages, but should not affect the top 1%. Acemoglu et al. (2020) study the impact of AI on labor markets, using establishment-level data on online vacancies with detailed occupational information in the US over the period from 2010 to 2018. They classify

establishments as “AI exposed” when their workers engage in tasks that are compatible with current AI capabilities. They document a rapid growth of vacancies for AI positions in AI-exposed establishments, especially after 2015. Moreover, AI exposure is associated with both a significant decline in some of the skills previously demanded in vacancies and the emergence of new skills, suggesting that AI is altering the task structure of jobs. However, they do not find an impact of AI exposure on employment or wages at the occupation or industry level, implying that AI is currently substituting for humans in a subset of tasks but it is not yet having detectable aggregate labor market consequences.

Similar conclusions about the effect of advanced ML technologies on the reorganization of tasks is reached by Brynjolfsson et al. (2018), who study the suitability of occupations for machine learning (SML), using the O*NET database. They find that ML will affect very different parts of the workforce than earlier waves of automation. In particular, they find that (i) most occupations in most industries have at least some tasks that are SML; (ii) few if any occupations have all tasks that are SML; and (iii) unleashing ML potential will require a significant redesign of the task content of jobs, as SML and non-SML tasks within occupations are unbundled and re-bundled. Therefore, the policymakers’ concern should be also directed towards the re-engineering of business processes. Along the same lines, Nedelkoska and Quintini (2018) highlight the significant changes that jobs will undergo as a result of the adoption of AI and ML. Their study estimates the risk of automation for jobs in 32 OECD countries using individual-level data on job tasks, finding that about 14% of jobs are highly automatable and that in 32% of jobs a significant share of tasks, but not all, could be automated, changing the skill requirements for these jobs.

The effect of AI technologies on firm labor productivity is the object of a recent study by Damoli et al. (2021). Using a unique database of 5257 AI patenting firms between 2000 and 2016, AI patent applications generate a significant effect on companies’ labor productivity, especially for small-medium enterprises and services industries. By doubling the number of AI patent applications, the predicted increase in labor productivity amounts to 3%. Notably, the study uses a comprehensive definition of AI that refers to the combination of software and hardware components including robotics. A key challenge in evaluating the role of new technologies in the labor market is the lack of micro-level information on technology adoption. As a matter of fact, much of the existing empirical work on the effect of AI in the labor market uses data on factory robotics (Acemoglu et al., 2020; Acemoglu & Restrepo, 2020a; Dauth et al., 2017; Koch et al., 2019), automation (Gregory et al., 2018), or ICT (Balsmeier & Woerter, 2019). A typical finding of this literature is that new technologies will especially impact mid-skilled, routine jobs. However, the empirical literature currently lacks sufficiently granular data on the type of technology adopted to test the conclusion of a rising gap between low- and high-skilled workers, and further research is needed on this issue.

2.2 | Innovation process and firm organization

AI can also play the role of human capital in the innovative production process, by changing the logic of discovery and the conduction of innovative activities. The role of AI within the innovation process is highlighted by Cockburn et al. (2019), who suggest that AI could represent a general-purpose “invention of a method of inventing”. Having collected and classified academic publications from 1955 to 2015, and patents from 1990 to 2014 as symbolic systems, learning systems, robotics, or “general” AI, they provide quantitative evidence on the evolution of these different areas, and document a meaningful shift in the application orientation of learning-oriented

publications, particularly after 2009. As a general-purpose invention of a method of invention, they argue that artificial intelligence technologies will likely affect also the organization of the innovation process. Policies that encourage transparency and sharing of core datasets may thus be critical tools for stimulating research productivity and innovation-oriented competition. More in general, policymakers' concern should be directed at the design of the incentives for the development and diffusion of these technologies and at ensuring that different potential innovators can gain access to these tools and to use them in a pro-competitive way. Because of the contribution of AI in the discovery process, it can also play an important role in science. Agrawal, McHale and Oettl (2019) focus on its role in supporting human researchers to improve the mechanism of discovery in scientific research. They formalize the idea that data-driven intelligence can solve problems that challenge human intelligence, in particular finding useful combinations in complex discovery spaces. In fact, the existing knowledge base interacts in highly complex ways and determines a massive number of potential new combinations, which must be searched and analyzed to discover those that provide valuable new knowledge. Meta technologies such as deep learning can aid the discovery process by allowing researchers to identify valuable combinations. By facilitating the access to data and knowledge, AI can improve prediction accuracy and discovery rates, thereby speeding up growth.

The effects of AI on the labor market might be more complicated when the firms' internal organization is accounted for. Technology may increase the complementarity between low-skilled and high-skilled workers, which increases the bargaining power of low-skilled workers (Aghion, Bergeaud, Blundell, & Griffith, 2019). In fact, the more innovative the firm, the more important it is to have high-ability low-occupation employees to make sure that the high-occupation employees within the firm can focus on the most difficult tasks (Garicano & Rossi-Hansberg, 2016), hence the need to select out those low-occupation employees which are not trustworthy. As a consequence, the prediction of a premium to skills may hold at the macroeconomic level, but perhaps it misses important aspects of firms' internal organization.

Low-skilled workers may benefit from AI also because they are those who ensure the "last mile" in AI production (Gray & Suri, 2017). Indeed, they perform important but low-profile human, micro-working, tasks in the "back-office" of AI, such as identifying objects on a photograph, adding labels to images, or correcting and sorting the data that help to train and test algorithms. Tubaro and Casilli (2019), by analyzing data from a detailed inventory of French-based micro-working platforms between June 2017 and March 2019 in the automotive industry, find that such a micro-work is a structural feature of today's AI production processes.

If low-skill workers can provide some necessary inputs to AI, the presence of high-skill workers might endogenously impact the innovativeness of the firm. In particular, the presence of workers executing abstract tasks, i.e. cognitive analytical and interpersonal activities, has a linear positive relationship with the propensity to innovate (Fonseca et al., 2019).

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, and move toward flatter organizational structures, thus speeding up the decentralization process caused by IT technologies (Bloom et al., 2014).

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 that 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) explains that it 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.

The empirical evidence on the impact of AI on labor, skills and firm organization is still scant but growing and it is reported in Table 1 where a comparison of these studies has been done in terms of data used in the analysis, the methodology employed and main findings.

Table 1 provides an overview of the studies described in this section relative to the effects of AI in the labor market.

3 | AI AND CONSUMERS

As user data is a fundamental input to AI systems, in this Section we analyze the effects of AI on consumers' behavior and surplus, focusing on the distortions that could be generated by the use of algorithms.

AI 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 function is particularly useful for consumers, especially because machines are more efficient and objective than human beings in selecting relevant and quality information, potentially leading to better matching and reduced search costs. In this respect, 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 & Elkin-Koren, 2017). In this way, algorithms help consumers to overcome behavioral biases and cognitive limits, make more rational choices and empower them against manipulative marketing techniques. Bundorf et al. (2019) run a randomized controlled trial in which they offer access to a decision-support tool incorporating algorithmic recommendations for choosing the cost-minimizing insurance plan. They find that the algorithmic advice significantly increases the probability to switch plan. Notably, however, the authors also find that the self-selection into software use is quantitatively important. In fact, many people who accept the algorithmic support were planning to switch the insurance plan in any case, whereas those who decline would most benefit from such decision-making support. This suggests that merely providing access to AI support is not sufficient to internalize its benefits.

The use of ML technologies may also present drawbacks from the consumers' point of view. ML technologies might produce selection biases, leading to a whole new range of policy concerns, especially given that the underlying predictive models are hardly interpretable or controllable by humans. Algorithmic unfairness can be originated either from biased algorithmic predictions or from biased algorithmic objectives (Cogwill & Tucker, 2020). The resulting biases, which are the object of the studies summarized in Table 2, may occur for two main reasons (Saurwein et al., 2015): first, they make predictions based on data that are endogenously generated; second, they incorporate the behavioral biases of human beings.

ML-based processes typically consider a large amount of data, including personal and demographic information. Since the learning processes of these algorithms are black boxes, they may lead to unintended discriminatory outcomes. For example, a heated debate sparked recently about

TABLE 1 Effects of AI on firms and labor

	Effect of AI	Methodology and findings
Bessen (2018)	Employment	<p>Model of AI as a labor-augmenting factor</p> <ul style="list-style-type: none"> The effect of AI on jobs depends on the elasticity of demand: new technologies replace labor with machines, but they also decrease prices, i.e. increase demand. If demand increases sufficiently, employment grows.
Acemoglu et al. (2020)	Employment	<p>AI substitutes human labour in a subset of tasks, altering the task structure of jobs</p> <ul style="list-style-type: none"> Establishment-level data on online vacancies in the US (2010-2018) Growth of vacancies for AI positions in AI-exposed establishments No significant impact of AI exposure on employment or wages at the occupation or industry level.
Agrawal, McHale and Oetl (2019)	Productivity in science	<p>Model where AI produces new knowledge</p> <ul style="list-style-type: none"> More accurate predictions can speed up growth via higher discovery rates.
Cockburn et al. (2019)	Innovation	<p>AI as a general-purpose invention of a method of invention</p> <ul style="list-style-type: none"> Evidence based on data on scientific publications (1955-2015) and patents (1990-2014) Shift in the importance of application-oriented learning research since 2009.
Brynjolfsson et al. (2018)	Firm organization	<p>AI as a substitute of human labor</p> <ul style="list-style-type: none"> O*NET data for 964 occupations in the US, joined to 18,156 tasks at occupation level Few jobs can be fully automated using ML ML potential can be exploited only after significant job redesign
Webb (2020)	Firm organization	<p>AI as a substitute of high-skills labor</p> <ul style="list-style-type: none"> O*NET data on occupations in the US and data on patent descriptions Workers in the 90th wage percentile are most exposed to AI AI might reduce the 90:10 wage inequality
Tubaro and Casilli (2019)	Firm organization	<p>Low profile micro-work is an input of AI</p> <ul style="list-style-type: none"> Data on 11 micro-working platforms in France (2017-2019) The development of AI solutions increases the relevance of micro-workers

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

TABLE 2 Effects of algorithms and AI on consumers and behavioral biases

	Effect of AI	Methodology and findings
Bundorf et al. (2019)	Reduction of search costs	Randomized, controlled trial of decision support software for choosing health insurance plans: <ul style="list-style-type: none">• Algorithmic expert recommendation significantly increases plan switching, cost savings, time spent choosing a plan, and choice process satisfaction• More “active shoppers” are more likely to use the decision-making support tool (evidence of self-selection)
Sweeney (2013)	Algorithmic bias (racial)	Distribution of ads by Google AdSense using a sample of racially associated names. Results suggest significant discrimination in ad delivery based on searches of 2184 racially associated personal names across two websites.
Angwin et al. (2016)	Algorithmic bias (racial)	Algorithm used in courtrooms for predicting recidivism misclassifies defendants in different ways: black defendants are often predicted to be at a higher risk of recidivism than they are; white defendants are predicted to be less risky than they are.
Miller and Tucker (2018)	Algorithmic bias (racial)	Empirical analysis of the efficacy of an algorithm that attempts to predict a person’s ‘ethnic affinity’ from their data online. The ad algorithm tends to overpredict the presence of African American in states where there is a historical record of discrimination against African Americans.
Datta et al. (2015)	Algorithmic bias (gender)	Browser-based experiments finding evidence of discrimination in the Ad Setting webpage
Lambrecht and Tucker (2019)	Algorithmic bias (gender)	Field test on ads for careers in the STEM sector. A cost optimizing, gender-neutral algorithm shows fewer ads to women relative to men

debate about the fairness criterion that should be used (Chouldechova, 2017). In this vein, Kosinski et al. (2013) report that someone liking (or disliking) ‘Curly Fries’ on Facebook is predictive of intelligence, hence it could be used as a screening device by algorithms whose goal is to identify desirable employees or students.

Discrimination might also arise from crowding-out effects. Lambrecht and Tucker (2019) show that an ad for jobs in the Science, Technology, Engineering and Math fields is less likely to be shown to women, even though 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 reacts 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 has a higher return on investment.

Algorithmic flaws might also originate from correlations in behavior. Miller and Tucker (2018) find that an advertising algorithm tends to over-predict 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 (Bertocchi & Dimico, 2014; Sokoloff & Engerman, 2000). In turn,

low-income people are more likely to use social media to express interest in celebrity 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 based on endogenously generated data (Tucker, 2019). Chander (2017) argues that the problem is not the black box of the algorithm, but the real world on which it operates. 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 the information according to the preferences of online users, leading to echo chambers (Sunstein, 2009; Claussen et al., 2019) and filter bubbles (Pariser, 2011). For example, ML algorithms implemented by search engines provide readers with news matching their own beliefs and preferences, but this effect depends on the amount of data the algorithm can use. Algorithmic recommendations on average receive more clicks than the human-curated control condition but only if the algorithm can use a relevant amount of individual-level data. This implies that a human editor is still better at identifying the taste of the average reader when an algorithm has limited data (Claussen et al., 2019). Product recommendations are also biased towards similar content to previous purchases. 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 (Vike-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).

AI outcomes may turn out discriminatory also because the algorithm itself will learn to be biased based on the behavioral data that feeds it. Documented alleged algorithmic biases span from 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 see 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 (2019) notes that people could myopically reveal sensitive information that could harm them in the future, a problem aggravated by some properties of data, like data persistence, data repurposing and data spillovers. 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 is created, there is uncertainty about how such data could be used in the future. Finally, there are also potential spillovers for others who did not provide the information, but are somehow affected by it.

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.

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

TABLE 3 Effects of algorithms and AI on markets and competition

	Application	Methodology and findings
Chen et al. (2016)	Pricing	Pricing algorithm in the form of a logit model for predicting the probability that a customer purchases a product at a given price.
Dubé and Misra (2017)	Pricing	Regression-based method for selecting the most “predictive” customer features, which capture the influence of price and demand, and contribute to customers’ price sensitivities.
Ezrachi and Stucke (2016)	Competition policy	Normative paper: computerised agents may be involved in anticompetitive collusion and antitrust policy challenges are discussed.
Klein (2019)	Competition policy	Simulations with pricing algorithms. Q-learning algorithms that compete sequentially learn to collude on prices.
Gautier et al. (2020)	Competition policy	Normative paper on algorithmic price discrimination and tacit collusion, discussed from an economic, technological and legal perspective.
Calvano et al. (2020)	Competition policy	Simulations with pricing algorithms. Q-learning algorithms that compete simultaneously learn to collude on prices.
Kosinski et al. (2013)	Consumer discrimination	Logistic/linear regression predicting individual psychodemographic profiles from Facebook likes. Facebook likes can be used to automatically and accurately estimate a wide range of personal private attributes

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 (Acquisti et al., 2015, 2016).

4 | AI, MARKETS AND COMPETITION

The exploitation of AI technologies has been described as a game changer (Ezrachi & Stucke, 2016) and it is expected to have a massive impact on existing markets. Some of these effects accrue to the use of digital technologies, which entail a reduction of search costs, replication costs, transportation costs, tracking costs and verification costs (Goldfarb & Tucker, 2019), although the effects are mediated by the individual’s personal characteristics (Castellacci & Tveito, 2018). Digital markets are at the forefront of competition policy and, in the last few years, antitrust authorities around the world have opened many investigations on digital platforms and issued or commissioned dozens of studies or expert reports that are focused on understanding the general competitive dynamics of markets such as online search, social media, e-commerce/marketplaces, and mobile operating systems (see Lancieri & Sakowski, 2021, for a survey). Despite its relevance, the goal of this Section is to focus on how the implementation of artificial intelligence and ML-based algorithms could affect market mechanisms and outcomes. We describe both the positive and negative effects of AI on competition and summarize them in Table 3.

4.1 | Pro-competitive effects of AI

The widespread use of ML is undoubtedly associated to significant efficiency effects, which benefit firms as well as consumers.

On the supply side, ML-based algorithms can promote static efficiency 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, ML is currently being employed by insurance companies to better assess the risk of customers, make automatic offers, and even process claims. The Economist (2017) reports that a policyholder can now receive the reimbursement three seconds after filing the claim on the app. In these three seconds, the machine can review the claim, run 18 anti-fraud algorithms, approve it, send payment instructions to the bank, and inform the customer. In the financial sector, ML is increasingly used to execute portfolio decisions, with AI systems choosing which stocks to buy and sell (The Economist, 2019). Supply-side efficiency improvements 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, rivals' prices and demand fluctuations. 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 et al., 2019; OECD, 2015). Indeed, ML-based algorithms have been used to develop new offerings, thus promoting market entry (OECD, 2016a, 2016b). For example, new “Intelligent Transport Systems” can be developed. These services are based on information and communication technologies and are applied to transport, including infrastructure and vehicles, traffic management, and interfaces between road and other modes of transport. Car-makers' business models already forecast a shift from selling cars and buses to selling ‘travel time well spent’ in which they collaborate with digitally savvy companies (OECD, 2016a). Within financial services, innovations like peer-to-peer lending involve reliable credit scoring systems, and some innovative players (like Alibaba in China and Upstart in the US) developed credit scoring mechanisms and income prediction models that grant a competitive advantage to their business model and that the banks are increasingly adopting (OECD, 2016b).

Furthermore, ML can promote competition by making information better organized and accessible for consumers. For instance, AI-powered search engines provide information on dimensions of competition other than prices, such as quality, to significantly reduce search and transaction costs and information asymmetries.⁶

AI might also play an important role in deterring collusion. Indeed, ML algorithms can help firms to better forecast demand and thus tailor prices to demand conditions. This implies that ML also increases each firm's temptation to deviate to a lower price in periods of high predicted demand. Hence, better forecasting and algorithms may lead to lower prices and increase consumer surplus as a consequence (Miklos-Thal & Tucker, 2019).

Moreover, algorithms can be usefully employed by Antitrust authorities as a detection tool to identify instances of coordination between suppliers and collusive pricing. Indeed, data-driven approaches have been proposed to detect bidding anomalies and suspicious bidding patterns across large data sets (OECD, 2017). For example, the Korea Fair Trade Commission has on several occasions succeeded in detecting bid-rigging conspiracies by screening procurement bidding

data. Akhgar et al. (2016) also suggest that ML-based algorithms could be applied to identify hidden relationships as an indicator of collusion in public tenders.

4.2 | Anti-competitive effects of AI

The adoption of AI technologies on a large scale is potentially associated with negative effects too, which could result in a reduction of the efficient functioning of the competitive mechanism. AI and ML are expected to exacerbate the typical market failures already highlighted for digital markets, caused by significant economies of scale, considerable network externalities and large switching costs on the demand side of the industries (Varian, 2019). In what follows we discuss the most immediate implications in terms of competition, which might call for the attention of policymakers, with a specific focus on the impact of algorithms on firms' incentives to collude.

Firms' pricing decisions are increasingly delegated to ML-based algorithms (Chen et al., 2016), which can account for 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 enhanced ability to recognize patterns within increasingly large datasets by ML algorithms helps a finer targeting and segmentation of the market than in the past (Milgrom & Tadelis, 2019). Better targeting dramatically enlarges the scope for price discrimination. First-degree price discrimination, so far only a theoretical possibility, could become a reality because of ML. Moreover, in an algorithm-driven environment, discrimination can be subtler than classical price discrimination, and take the form of behavioral discrimination (Ezrachi & 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 policymakers to uncover such practices, few instances of first-degree price discrimination have 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 profit increases 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 an 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. (2020) 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 concerns may limit the use of price discrimination by firms, 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).

Since data is an essential input for the algorithm, the control over consumers' personal information not only helps to construct a more efficient algorithm but is also the key element for market control (Cremer et al., 2019). As reported by The Economist (2017)⁷, data is the “world's most

valuable resources” and their exploitation is at the core of the new business models of digital platforms and their algorithms. Moreover, data is nonrival, and leads to potentially large gains when it is broadly used. In this setting, the ownership of the data affects both consumers’ and firms’ behavior. When data property rights are assigned to consumers, they will optimally balance their concerns for privacy against the economic gains from selling the data to all interested parties (Jones & Tonetti, 2020). On the other side, if the ownership of large datasets is in the firms’ hands, it 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, 2019).

The recent literature on data economics has emphasized how access to customer-level data may provide to firms private information, which could be used for a competitive advantage (Casadesus-Masanell & Hervas-Drane, 2015; Montes et al., 2019). In particular, information selling allows firms not only to extract surplus from consumers but also to increase competition since firms will then set their prices more aggressively. From the policy perspective, the problem that comes out is the incentive of the data broker to share data. Indeed, the data broker will prefer to sell information to only one of the competitors to soften competition and then extract the monopolistic rent through data selling. In other words, the goal of the data broker is to limit market competition to increase the value of information (Montes et al., 2019).

Data-enabled learning however may take many forms, as suggested by Hagiu and Wright (2020). They distinguish across-users learning (i.e., more users generate more data) and within-users learning (i.e., higher usage intensity generates more individual data) that creates endogenous switching costs to consumers and provides a competitive advantage to the incumbent firms. In this setting, imposing data sharing may induce firms to compete less aggressively for data acquisition, implying to pay a lower price to consumers for their data, thus potentially lowering consumer surplus. By studying the interaction between these two types of learning, Schafer and Sapi (2020) provide evidence supporting the claim that data as an input into machine learning constitutes a source of market power. They find that a search engine with access to longer user histories may improve the quality of its search results faster than an otherwise equally efficient rival with the same size of user base but access to shorter user histories. The sharing of consumers’ data generates a negative externality that reduces consumers’ surplus, owing to their loss of privacy. This externality might be corrected by allowing consumers to be compensated for their data. However, imposing such a price on data might also have negative side effects. Acemoglu et al. (2019) show that the price of data is affected by data externalities and might lead to excessive data sharing. Moreover, when the data provided by one consumer has a negative externality on others, the price of data can be substantially below the value of information to the platform (Bergemann et al., 2019).

From a policymaking point of view, the externalities arising from data sharing and use call for some sort of data regulation. Indeed, the massive and unprecedented scale of data is creating serious concerns by policymakers and the public for their impact on market competition and the large loss in terms of privacy.

Competition commissions throughout the world are expressing concerns about the implications of data control for competition, consumers, and society. For example, a report for the European Commission (Cremer et al., 2019) and another one for the US (Stigler Center, 2019) point out the risks for competition in the market of the new digital giants. Accordingly, they both invoke a specific extension of antitrust rules regarding rules for structural separation, data access and sharing as well as the creation of a new Authority for data regulation. The Australian Competition and Consumer Commission also observed that the breadth and scale of the user data collected by

platforms are relevant both for the assessment of their market power and for consumer concerns (ACCC, 2019).

Facing the challenges of digitalization might require a revision of the current regulatory framework, and indeed many countries are considering policy changes in this area. For example, the UK government is currently establishing a new regulatory framework for digital markets, whereby companies that allow users to share user-generated content will be subject to an independent regulator⁸. To this end, the UK Competition and Market Authority recommends that the government establishes a pro-competition regulatory regime for online platforms “to enforce a code of conduct to govern the behavior of platforms with market power, ensuring concerns can be dealt with swiftly, before irrevocable harm to competition can occur” (Competition and Markets Authority, 2020).

Another serious concern is about the potential role of data-enabled algorithms in facilitating collusion. Algorithmic pricing may facilitate collusion via two main channels. First, ML algorithms learn to react to rivals’ prices much more quickly than human beings (Ezrachi & Stucke, 2016; Mehra, 2015). Because of the frequent interactions, defection from a collusive agreement is punished more promptly and gains from defection are reaped for a shorter time. Thus, automating a firm’s price response to rival’s prices through an algorithm provides a firm with an advantage relative to its peers in terms of frequency of price changes and leads to higher prices relative to the competitive ones (Brown & MacKay, 2019).

Second, ML-based algorithms actively learn the optimal strategy purely by trial and error, by intentionally experimenting sub-optimal prices. These kinds 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 (2019) shows that simple algorithmic agents could learn to collude in a sequential move game. A similar finding is obtained by Calvano et al. (2020) even when moves are simultaneous. They run an experiment with AI-powered pricing algorithms which interact in a controlled environment of computer simulations. The study finds that AI pricing agents systematically learn to play sophisticated collusive strategies without communicating with one another. They charge supra-competitive prices and mete out punishments that are larger, the larger the deviation and are finite in duration, with a gradual return to pre-deviation prices. Differently from collusion between human subjects, collusive strategies played by AI agents are robust to perturbations of cost or demand, number of players, asymmetries and forms of uncertainty.

The overall impact of algorithm pricing is thus not conclusive, with some studies presenting potential positive effects in terms of lower prices (Miklos-Tal and Tucker, 2019) or the opposite, that is, higher prices with a high risk of collusion (Calvano et al., 2020). The empirical evidence is scant, but some results are already available. A recent analysis by Assad et al. (2020) on the use of algorithmic-pricing software in Germany’s retail gasoline markets shows that in duopolistic markets where both gas stations use an algorithm pricing market-level margins, prices increase by 28%. Overall, this result implies that the adoption of algorithmic pricing has affected competition facilitating tacit collusion in the German retail gasoline market.

From a policy standpoint, as algorithmic systems become more sophisticated, they are often less transparent, and it is more challenging to identify when they cause harm (Competition and Markets Authority, 2021). Moreover, not only detection of competitive harms, but also enforcement of Antitrust law becomes more challenging. From an enforcement point of view, a critical problem is that pricing through ML algorithms leaves no clear 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 especially in the presence of mass adoption of algorithmic pricing software. For example, US agencies require evidence of communication between the parties to determine that an agreement exists, and this may not be easy to establish where AI systems are concerned (Rab, 2019). 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 AI system, on the individual who used it or on the person (or entity) who benefited from the decision made by the system, even if consumers' harm was not consciously done?

The answers to such questions are not clear-cut at the moment, and even the realism of collusion by ML algorithms is presently an object of debate, given that real antitrust cases have not yet emerged. Gautier et al. (2020) observe that such a scenario might never materialize, as there are technical and market barriers that hinder the emergence of algorithmic tacit collusion outside the realm of lab experiments.

5 | CONCLUSIONS AND FUTURE RESEARCH

In this paper, we provide an overview of the many and multi-faceted economic effects of the recent technological advances in Artificial Intelligence that involve machine learning applications, drawing attention to those issues with the most urgent policy implications. We examine the effects of AI in the labor market, focusing on its implications on productivity, employment, firm organization and innovation process. Then we examine how AI contributes to shaping consumer behavior and market competition, by exploiting newly accessible data sources, data-enabled learning and preexisting behavioral biases of human beings.

The effects of AI in terms of labor market outcomes have largely dominated the policy discussion in recent years, with economists highlighting its challenges in terms of wage inequality and unemployment. Indeed, there is now a rising call for policies ranging from changes in patent length and capital taxation (Korinek & Stiglitz, 2019), to employment subsidies (Eissa & Liebman, 1996; Hotz et al., 2006), up to Universal Basic Income policies. Despite the policymakers' concern on the disruptive effects of AI in the labor market, few studies possess a sufficient granularity of data on the technology adopted at the firm-level to assess the extent of this impact. Indeed, most of the available empirical literature on the effect of AI on the labor market uses data on factory robotics and automation. Robotics often employs AI for processing data, but its economic use is quite specific, and centers on the automation of narrow tasks, that is, substituting machines for certain physical activities previously performed by humans (Acemoglu & Restrepo, 2020b). Conversely, the literature we have surveyed suggests that AI is a more pervasive technology, which includes various areas of research and poses different challenges depending on the production process and on the specificity of the industries where it is implemented. As such, the study of the effects of AI requires a much broader focus than just robotics, and has to take a further step to open up the black box of ML-related AI. Therefore, it should disentangle its effects on employment based on its specific applications in manufacturing as well as in the service industry, particularly in finance, banking, retailing and health care where the demand for its services is expected to grow significantly over the next few years.

The economic effects of data-enabled learning go beyond their impact on the labor market. As shown in the second part of this study, AI technologies can provide important and direct consumer benefits, through higher-quality and more accessible information. They also have a massive impact on the functioning of existing markets, on their boundaries and on ways with which firms interact between themselves and with consumers. This, however, may imply new threats for consumers' welfare, increasing the risk of new elusive forms of collusion and firms' exploitative practices.

Turning to a different perspective, state-of-the-art research suggests important implications on competition, the product market, and consumers. On the one hand, ML-based AI is expected to facilitate the entry of new firms, thus increasing competition, on the other hand it strengthens the market power of big tech companies. Which effect prevails is an empirical matter that deserves to be further explored. Moreover, AI is likely to influence the degree of vertical integration of digital markets. For example, Google is acquiring hundreds of startups developing AI solutions. The effects of mergers in digital markets have been recently studied by Gautier and Lamesch (2020), but the impact of mergers on AI developers is still to be addressed. More generally, evidence on collusive practices or real antitrust cases is still missing to support the theoretical predictions, thereby calling for further study.

Finally, the increasing pervasiveness of computers calls for an understanding of how humans actually behave in interaction with intelligent machines. One important problem is caused by people's irrational behavior, which not only leaves room for the exploitation of users' data and entails privacy losses, but also originates biases that could be amplified by algorithms.

Policymakers will face unprecedented challenges to face the new complex and rapidly evolving environment and to fill the gap between policy and enforcement concerning the ability to find evidence of human involvement where machines or algorithms indeed facilitate anti-competitive behavior (Rab, 2019). 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 occur, 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 pricing process adopted by firms, thereby making it challenging for Antitrust authorities to detect and punish anticompetitive conduct. Indeed, although AI-based tools may provide policymakers with precious support to and improve policy accuracy, there are limits to the scope of their action and, more importantly, AI "does not help us balance interests or engage in politics" (Goolsbee, 2018, p. 8). 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, as the control over search results can also be exploited for political interests (Epstein & Robertson, 2015). While extremely relevant, the issue of data agglomeration and exploitation for political purposes is beyond the scope of this survey.

ACKNOWLEDGMENTS

This work has been partially supported by "Ministero dell'Istruzione, dell'Universita e della Ricerca" Award "TESUN- 83486178370409 finanziamento dipartimenti di eccellenza CAP. 1694 TIT. 232 ART. 6".

Open access funding enabled and organized by Projekt DEAL.

ORCID

Laura Abrardi  <https://orcid.org/0000-0002-3910-7097>

Carlo Cambini  <https://orcid.org/0000-0002-7471-8133>

Laura Rondi  <https://orcid.org/0000-0002-7683-1164>

ENDNOTES

- ¹ Accessible at <https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449>
- ² In this respect, (Agrawal et al., 2019a) note that automation is just one of the potential consequences of AI.
- ³ In its White Paper presented on February 19, 2020, the European Commission (2020) envisages a regulatory framework for Artificial Intelligence where rules should be applied to address the risks associated with AI applications, to guarantee consumer protection, fair commercial practices and protection of personal data and privacy.
- ⁴ The potential of intelligent machines of substituting human labor is the focus of studies that adopt a broader definition of AI, more closely related to the concept of automation, and model it within a capital-labor productivity framework. In these studies, machines typically substitute workers either by increasing the return of capital (Nordhaus, 2020), by providing a new factor of production, “robotic labor” (DeCanio, 2016), or by expanding the set of tasks produced by machines (Acemoglu & Restrepo, 2018b).
- ⁵ In particular, clinical laboratory technicians, chemical engineers, optometrists, power plant operators and dispatchers.
- ⁶ ML-based algorithms, for example, are necessary to analyze the quality-related information contained in text data (Gentzkow et al., 2019, provide a survey of the relevant statistical methods and applications to analyze text).
- ⁷ The Economist, May 6, 2017, “Regulating the internet giants. The world’s most valuable resource is no longer oil, but data”.
- ⁸ Online Harms White Paper, Presented to Parliament on April 2019 and available at www.gov.uk/government/publications.

REFERENCES

- ACCC. (2019). Digital platforms inquiry, final report. Report 06/19 1545, *Australian Competition and Consumer Commission*.
- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2020). AI and jobs: Evidence from online vacancies. NBER Working Paper 28257.
- Acemoglu, D., Lelarge, C., & Restrepo, P. (2020). Competing with robots: Firm level evidence from France. *AEA papers and proceedings*, 110, 383–88.
- Acemoglu, D., Makhdoumi, A., Malekian, A., & Ozdaglar, A. (2019). Too much data: Prices and inefficiencies in data markets. *NBER Working Paper no. 26296*.
- Acemoglu, D., & 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.
- Acemoglu, D., & Restrepo, P. (2018b). Modeling automation. *AEA Papers and Proceedings*, 108, 48–53.
- Acemoglu, D., & Restrepo, P. (2019). Artificial intelligence, automation and work. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Acemoglu, D., & Restrepo, P. (2020a). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244.
- Acemoglu, D., & Restrepo, P. (2020b). The wrong kind of AI? Artificial Intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society, Cambridge Political Economy Society*, 13(1), 25–35.
- Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509–514.
- Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492.
- Akerman, A., Gaarder, I., & Mogstad, M. (2015). The skill complementarity of broadband internet. *The Quarterly Journal of Economics*, 130(4), 1781–1824.
- Aghion, P., Bergeaud, A., Blundell, R. & Griffith, R. (2019). The innovation premium to soft skills in low-skilled occupations. CEP Discussion Papers dp1665, Centre for Economic Performance, LSE.

- Aghion, P., Jones, B., & Jones, C. (2019). Artificial intelligence and economic growth. In A. Agrawal, Gans & Goldfarb, A. (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Agrawal, A., Gans, J., & Goldfarb, A. (2016). The simple economics of machine intelligence. *Harvard Business Review*, 17.
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). Human judgement and AI pricing. *AEA Papers and Proceedings*, 108, 58–63.
- Agrawal, A., Gans, J., & Goldfarb, A. (2019a). *The economics of artificial intelligence: An agenda*. Chicago and London: University of Chicago Press.
- Agrawal, A., Gans, J., & Goldfarb, A. (2019b). Economic policy for artificial intelligence. *Innovation Policy and the Economy*, 19, 139–159.
- Agrawal, A., McHale, J., & Oettl, A. (2019). Finding needles in haystacks: Artificial intelligence and recombinant growth. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Akhgar, B., Bayerl, P. S., & Sampson, F. (2016). *Open source intelligence investigation: From strategy to implementation*. Springer International Publishing.
- 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).
- Assad, S., Clark, R., Ershov, D., & Xu, L. (2020). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *CESifo Working Paper No. 8521*.
- Athey, S., Calvano, E., & Gans, J. S. (2018). The impact of consumer multi-homing on advertising markets and media competition. *Management Science*, 64(4), 1574–1590.
- Athey, S., Bryan, K. A., & Gans, J. S. (2020). The allocation of decision authority to human and artificial intelligence. *AEA Papers and Proceedings*, 110, 80–84.
- Autor, D. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30.
- Autor, D., & Salomons, A. (2018). Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity*, 1–63.
- Balsmeier, B., & Woerter, M. (2019). Is this time different? How digitalization influences job creation and destruction. *Research Policy*, 48(8), 103765.
- Bar-Ilan, J. (2007). Google bombing from a time perspective. *Journal of Computer-Mediated Communication*, 12(3), 910–938.
- Belleflamme, P., Lam, W., & Vergote, W. (2017). Price discrimination and dispersion under asymmetric profiling of customers, CORE discussion paper, Louvain-la-Neuve, Belgium.
- Bergemann, D., Bonatti, A., & Gan, T. (2019). The economics of social data. *Cowles Foundation Discussion Paper 2203*.
- Bessen, J. (2018). AI and jobs: The role of demand. *NBER Working Paper no. 24235*
- Bessen, J. (2018). AI and jobs: The role of demand. *NBER Working Paper no. 24235*
- Blake, T., Nosko, C., & Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1), 155–174.
- Bloom, N., Garicano, L., Sadun, R., & Van Reenen, J. (2014). The distinct effects of information technology and communication technology on firm organization. *Management Science*, 60(12), 2859–2885.
- Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103, 347–356.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford: Oxford University Press.
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies 'Engines of growth'? *Journal of Econometrics*, 65(1), 83–108.
- Brown, Z., & MacKay, A. (2019). Competition in pricing algorithms. Working paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3485024
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. New York, NY: W.W. Norton, 2014.
- 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.

- Bundorf, K., Polyakova, M., & Tai-Seale, M. (2019). How do humans interact with algorithms? Experimental evidence from health insurance. *NBER Working Paper no. 25976*.
- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing and collusion. *American Economic Review*, *110*(10), 3267–3297.
- Casadesus-Masanell, R., & Hervas-Drane, A. (2015). Competing with privacy. *Management Science*, *61*(1), 229–246.
- Castellacci, F., & Tveito, V. (2018). Internet use and well-being: A survey and a theoretical framework. *Research Policy*, *47*, 308–325.
- Chander, A. (2017). The racist algorithm? *Michigan Law Review*, *115*(6), 1023.
- Chen, L., Mislove, A., & 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.
- Chevalier, J. (2019). Antitrust and artificial intelligence: Discussion of Varian. In A. Agrawal, Gans & Goldfarb, A. (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, *5*(2), 153–163
- Claussen, J., Peukert, C., & Sen, A. (2019). The Editor vs. the Algorithm: Targeting, Data and Externalities in Online News. <https://ssrn.com/abstract=3399947> or <https://doi.org/10.2139/ssrn.3399947>
- Cockburn, I., Henderson, R., & Stern, S. (2019). The impact of artificial intelligence on innovation: An exploratory analysis. In A. K. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: an Agenda*. University of Chicago Press.
- Cogwill, & Tucker, C. (2020). Economics, fairness and algorithmic bias. *The Journal of Economic Perspectives*, forthcoming.
- Competition and Markets Authority. (2020). Online platforms and digital advertising. *Market study final report*.
- Competition and Markets Authority. (2021). Algorithms: How they can reduce competition and harm consumers. *Research and analysis final report*.
- 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*.
- Damioli, G., Van Roy, V., & Vertesy, D. (2021). The impact of artificial intelligence on labor productivity. *Eurasian Business Review*, *11*, 1–25.
- Datta, A., Tschantz, M. C., & Datta, A. (2015). Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies*, (1), 92–112.
- Dauth, W., Findeisen, S., Südekum, J., & Wößner, N. (2017). German robots –The impact of industrial robots on workers. *CEPR Discussion Papers 12306*.
- DeCanio, S. J. (2016). Robots and humans – complements or substitutes? *Journal of Macroeconomics*, *49*, 280–291.
- Dubé, J.-P., & Misra, S. (2017). Scalable price targeting. *NBER Working Paper no. 23775*
- Eissa, N., & Liebman, J. B. (1996). Labor supply response to the earned income tax credit. *Quarterly Journal of Economics*, *111*(2), 605–637.
- Epstein, R., & 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).
- European Commission. (2020). White Paper on Artificial Intelligence: A European approach to excellence and trust. February 19, 2020. <https://ec.europa.eu/info/files/white-paper-artificial-intelligence-european-approach-excellence-and-trust>
- Ezrachi, A., & Stucke, M. E. (2016). Virtual competition. *Journal of European Competition Law & Practice*, *7*(9), 585–586.
- Gal, M. S., & Elkin-Koren, N. (2017). Algorithmic consumers. *Harvard Journal of Law and Technology*, *30*(2), 309–352.
- Gal, M. S., & Elkin-Koren, N. (2017). Algorithmic consumers. *Harvard Journal of Law and Technology*, *30*(2), 309–352.
- Garicano, L., & Rossi-Hansberg, E. (2016). Organization and inequality in a knowledge economy. *The Quarterly Journal of Economics*, *121*(4), 1383–1435.
- Gautier, A., Ittoo, A., & Van Cleynenbreugel, P. (2020). AI algorithms, price discrimination and collusion: A technological, economic and legal perspective. *European Journal of Law and Economics*, *50*(3), 405–435.
- Gautier, A., & Lamesch, J. (2020). Mergers in the digital economy. *CESifo Working Paper No. 8056*.

- Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3), 535–574.
- Goldfarb, A., & Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1), 3–43.
- Goolsbee, A. (2018). Public policy in an AI economy. *NBER Working Paper 24653*.
- Graetz, G., & Michaels. (2018). Robots at work. *Review of Economics and Statistics*, 100(5), 753–768.
- Gray, M., & Suri, S. (2017). The humans working behind the AI curtain. *Harvard Business Review*, (9), 2–5.
- Gregory, T., Salomons, A., & Zierahn, U. (2018). Racing with or against the machine? *Evidence from Europe*, CESifo Working Papers n. 7247.
- Hagiu, A., & Wright, J. (2020). Data-enabled learning, network effects and competitive advantage. Working paper.
- Hotz, V. J., Mullin, C. H., & Scholz, J. K. (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.
- Jin, G. Z. (2018). Artificial Intelligence and consumer privacy. *NBER Working Paper No. 24253*.
- Jones, C. I., & Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9), 2819–2858.
- Kaplan, J. (2016). *Artificial Intelligence: What everyone needs to know*. Oxford University Press.
- Klein, T. (2019). 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.
- Koch, M., Manuylov, I., & Smolka, M. (2019). Robots and firms. Economics Working Papers 2019-05. *Department of Economics and Business Economics*. Aarhus University.
- Korinek, A., & Stiglitz, J. E. (2019). Artificial Intelligence and its implications for income distribution and unemployment. In A. Agrawal, Gans & Goldfarb, A. (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15), 5802–5805.
- Kotlikoff, L., & Sachs, J. D. (2012). Smart machines and long-term misery, *NBER Working Paper No. 18629*.
- Lambrecht, A., & Tucker, C. (2019). Algorithmic bias? An empirical study into apparent gender-based discrimination in the display of STEM career ads. *Management Science*, 65(7), 2966–2981.
- Lancieri, F., & Sakowski, P. M. (2021). Competition in digital markets: A review of expert reports. *Stanford Journal of Law, Business & Finance*, 26(1), 65-170
- Lu, Y., & Zhou, Y. (2021). A review on the economics of artificial intelligence. *Journal of Economic Surveys*, forthcoming.
- Makridakis, S. (2017). The forthcoming Artificial Intelligence (AI) revolution: Its impact on society and firms. *Futures*, 90, 46–60.
- Mehra, S. K. (2015). Antitrust and the Robo-Seller: Competition in the time of algorithms, Minnesota. *Law Review*, 100.
- Miklos-Thal, J., & Tucker, C. (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science*, 65(4), 1552–1561.
- Milgrom, P. R., & Tadelis, S. (2019). How Artificial Intelligence and machine learning can impact market design. NBER Chapters, In Agrawal, A., Gans, J. & Goldfarb, A., *The economics of artificial intelligence: An agenda*, pages 567–585, University of Chicago Press.
- Miller, A., & Tucker, C. (2018). Historic patterns of racial oppression and algorithms. *Mimeo*, MIT.
- Mitchell, T., & Brynjolfsson, E. (2017). Track how technology is transforming work. *Nature*, 544(7650), 290–292.
- Montes, R., Sand-Zantman, W., & Valletti, T. (2019). The value of personal information in online markets with endogenous privacy. *Management Science*, 65(3), 955–1453.
- Nedelkoska, L., & Quintini, G. (2018). Automation, skills use and training. *OECD Social, Employment and Migration Working Papers, No. 202*. Paris: OECD Publishing.
- Nordhaus, W. (2020). Are we approaching an economic singularity? Information technology and the future of economic growth. *American Economic Journal: Macroeconomics*. Forthcoming
- OECD. (2015). *Data-Driven Innovation: Big Data for Growth and Well-Being*. Paris: OECD Publishing. <https://doi.org/10.1787/9789264229358-en>.
- OECD (2016a). 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).
- OECD. (2016b). 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).

- OECD. (2017). Algorithms and Collusion: Competition Policy in the Digital Age www.oecd.org/competition/algorithms-collusion-competition-policy-in-the-digital-age.htm
- Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. New York, NY: Penguin.
- Petit, N. (2017). Antitrust and Artificial Intelligence: A research agenda. *Journal of European Competition Law & Practice*, 8(6).
- Rab, S. (2019). Artificial intelligence, algorithms and antitrust. *Competition law journal*, 18(4), 141–150.
- Saurwein, F., Just, N., & Latzer, M. (2015). Governance of algorithms: Options and limitations. *Info*, 17(6), 35–49.
- Schafer, M., & Sapi, G. (2020). Learning from data and network effects: The example of internet search. *DIW Berlin Discussion Paper No. 1894*.
- Shiller, B. R. (2014). First degree price discrimination using big data. *Economics Department*. Brandeis University, MA.
- Sokoloff, K. L., & Engerman, S. L. (2000). Institutions, factor endowments, and paths of development in the new world. *Journal of Economic Perspectives*, 14(3), 217–232.
- Stigler Center. (2019). Stigler Committee on Digital Platforms, Final Report. *Stigler Center for the Study of the Economy and the State*, The University of Chicago Booth School of Business. <https://research.chicagobooth.edu/stigler/media/news/committee-on-digital-platforms-final-report>.
- Sunstein, C. (2002). *Republic.com*, Princeton, NJ: Princeton University Press.
- Sunstein, C. (2009). *Republic.com 2.0*, Princeton, NJ: Princeton University Press.
- Sweeney, L. (2013). Discrimination in online Ad delivery. *Communications of the ACM*, 56(5), 44–54.
- Taddy, M. (2019). The technological elements of artificial intelligence. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Tirole, J. (2017). *Economics for the common good*. Princeton University Press, 2017.
- 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.
- The Economist (2019). The stockmarket is now run by computers, algorithms and passive managers, www.economist.com/briefing/2019/10/05/the-stockmarket-is-now-run-by-computers-algorithms-and-passive-managers. Last accessed October 10th, 2019.
- Townley, C., Morrison, E., & Yeung, K. (2017). Big data and personalized price discrimination in EU competition law. *Yearbook of European Law*, 36, 683–748
- Tubaro, P., & Casilli, A. (2019). Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics*, 46, 333–345.
- Tucker, C. (2019). Privacy, algorithms and artificial intelligence. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Webb, M. (2020). The impact of artificial intelligence on the labor market. *Stanford University Working Paper*.
- Varian, H. (2019). Artificial intelligence, economics, and industrial organization. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Vike-Freiberga, V., Däubler-Gmelin, H., Hammersley, B. & Pessoa Maduro, L. M. P. (2013). A free and pluralistic media to sustain European democracy. Report of High Level Group on Media Freedom and Pluralism.

How to cite this article: Abrardi L, Cambini C, Rondi, L. Artificial intelligence, firms and consumer behavior: A survey. *Journal of Economic Surveys*. 2021;1–23. <https://doi.org/10.1111/joes.12455>