Draft

Artificial intelligence and markets

Martti Hetemäki
February 2018

Abstract: The main point of the paper is that AI is likely to affect labor market and rest of the economy via changes in market efficiency. Such an approach enables a testable analysis of AI's economic effects. Moreover, it provides a basis to identify potential policies and the time line with which AI affects the economy.

The paper examines two effects of AI. The first effect is based on a hypothesis that AI will, in general, improve efficiency of markets. This holds if, in general, AI is required to reach rational decisions based on available data. The increased market efficiency will aggravate effects of existing economic distortions.

The second effect is based on a hypothesis that AI may induce new distortions. Research on algorithmic high-frequency trading indicates that AI can give rise to new distortions which are neither easy to detect nor self-correcting. This effect may be amplified by AI's efficiency improving general effect.

The paper puts forward two propositions on how AI affects employment via the market effects.

The paper argues that AI affects the economy, via the markets, earlier than it automates jobs on a larger scale. AI's potential economic effects and their likely speed stress the need for proactive policies.

The paper argues that required policies would amount to addressing existing economic distortions and to boosting skills. These policies would be warranted even in the absence of AI. Moreover, they are likely to be policies that would facilitate sustained job creation in the face of AI's technological effects via automation.

1 Introduction

In addition to the fundamental ethical and governance questions on artificial intelligence (AI), much of the debate on it has focused on AI's effect on employment. While this is a very important issue, economic theory may have rather little to say about it. For a survey on expert and non-expert views on technological unemployment due to AI, see Walsh (2017a).

Regarding the level of technology to which AI can lead, Walsh (2016) provides many reasons why AI is unlikely to lead to so called singularity. In this context singularity can be said to be a stage at which a machine would have sufficient intelligence to be able to redesign itself to improve intelligence. At that stage AI would basically replace human intelligence.

While singularity is unlikely, Walsh (2017b) sees a fast growth in AI given four exponential developments:

- Exponential growth in computing capacity
- Exponential improvement in accuracy of AI algorithms (e.g. reduced error rate in face recognition)
- Exponential growth of data
- Exponential growth in investment in AI

There are by now many surveys and reports on AI developments and policies, see, e.g., the reports by Stanford University (2016) and US government (2016).

This paper confines itself to examining effects of AI on the economy via markets. The main point of the paper is that AI is likely to affect the economy via changes in market efficiency. This provides a basis to identify both potential policies and the time line with which AI impacts the economy.

The paper examines two effects of AI. The first effect is based on a hypothesis that AI will, in general, improve efficiency of markets. This holds if, in general, AI is required to reach rational decisions based on all available data. The increased market efficiency will aggravate effects of existing economic distortions.

The second effect of AI is based on a hypothesis that AI may induce new distortions. Recent research on high frequency trading indicates that AI can give rise to new distortions which are neither easy to detect nor self-correcting. The second effect of AI may be amplified by AI's efficiency improving general effect.

The paper puts forward two propositions on how AI affects employment via the market effects. One proposition considers the effects of AI through more efficient product and capital markets under unchanged asymmetric information on labor market entrants. AI is unlikely to solve the informational problems for labor market entrants. The of lack of job market data on new labor market entrants in combination with more efficient product and capital markets impairs labor market functioning for those persons.

The more impaired labor market may arise also from the signal information on long-term unemployed persons (e.g. Kroft et al, 2013) when other markets function more efficiently. The reason for this impairment is the second best setting in which improved product and capital market efficiency aggravates the effects of existing labor market distortions.

The other proposition consider the effects of AI through changes in the skill requirements from cognitive to more non-cognitive skills making educational attainment a less valuable signal of skills. Technology and AI in particular are likely decrease the role of cognitive, and increase the role of non-cognitive, skills in the labor market. AI adds labor markets distortions by increasing asymmetric information when AI increases the role of non-cognitive skills in relation to cognitive skills resulting in a reduced signaling value of educational attainment. There is already empirical evidence from Finland and Sweden on increased role of non-cognitive skills (Jokela et al, 2017 and Edin et al, 2017).

The increased price information through internet has presumably already increased efficiency of product markets. ICT has also facilitated technology transfer and capital mobility and thus increased the effects of existing labor market imperfections, including those due to informational problems. This, in turn, may explain the observed reduced labor market fluidity (e.g., Davis and Haltiwanger, 2014). The increasing market efficiency due to AI can be expected to further hamper the functioning of certain segments of labor markets which may lead to further displacement of workers.

The paper argues that AI impacts the economy, via the markets, earlier than it automates jobs on a larger scale. The likely speed of AI's economic effects is a challenge for policy making. From the analysis one can draw policy conclusions. A separate paper analyses policies which would alleviate the adverse effects of AI.

The paper is organized as follows. Section 2 argues that AI will impact the economy via markets earlier than it automates jobs on a larger scale. Section 3 puts forward the proposition that AI will, in general, improve market efficiency. Section 4 presents a proposition that more efficient product and capital markets aggravate the effects of existing labor market distortions. Section 5 puts forward a proposition that AI adds labor markets distortions by increasing asymmetric information when AI increases the role of non-cognitive skills in relation to cognitive skills resulting in a reduced signaling value of educational attainment. Section 6 discusses the realism, practical relevance and implications of the results.

2 On the time line at which AI will impact the economy

On the speed or time line of technological unemployment due to AI, the before-mentioned survey on expert and non-expert views by Walsh (2017a) reaches the following conclusions:

> *"Whilst the experts predicted a significant number of occupations were at risk of automation in the next two decades, they were more cautious than people outside the field in predicting occupations at risk. Their predictions were consistent with their estimates for when computers might be expected to reach human level performance across a wide range of skills. These estimates were typically decades later than those of the non-experts. Technological barriers may therefore provide society with more time to prepare for an automated future than the public fear. In addition, public expectations may need to be dampened about the speed of progress to be expected in Robotics and AI."*

While in the survey by Walsh both groups of experts (experts on robotics and AI) predict that a large fraction of occupations were at risk of automation, the time line extends to the next couple of decades. In the following, it is argued that AI will impact the economy via markets faster than through automation.

To assess the speed at which AI enters the markets, one should first recall the four exponentials by Walsh (2017b) listed in the introduction. To gain further insight on the time line, one may ask two questions:

1. What are the factors affecting the use AI in the various markets?
2. What are the likely constraints on those factors?

The answer to the first question seems to be computer power, algorithms and data. Computer power does not seem to be a constraining factor. Algorithms appears also not to be a constraint since, in addition to academia, also firms appear ready to share their share. Apparently, e.g., open source software development and sharing the knowledge is in the industry as a whole. However, firms, such as Google, Amazon or Facebook, seem to be much less willing to share the data which they often obtain as a by-product of their customers using the services. Hence, it has been said, data is the oil or currency of AI.

Given that data, not the capacity to process it, is likely to be the limiting constraint to machine learning, one has to look at the development of available data to make predictions of the use of AI in markets.

While firms like Google, Amazon and Facebook are unlikely to share data to which they obtain from computer search and other uses of their customers, they will, of course, not have a monopoly to data on most markets. In fact, for many markets there is already today abundance of available data.

The above brief observations support the view that AI will impact the economy via markets much faster than it will impact the economy via automation of current jobs on a larger scale. Obviously, a much more systematic analysis of the available data for AI use by participants in different market would be warranted.

3 The effect of AI on market efficiency

In an analysis of mechanism design, Varian (1995) made the following observation some 20 years ago:

> *"Game theory has been justly criticized for its "hyper-rational" view of human behavior. However, such hyper-rationality may actually be an appropriate model for software agents: presumably software agents have much better computational powers than human beings. The whole framework of game theory and mechanism design may well find its most exciting and practical application with computerized agents rather than human agents, a point recognized by Rosenschein and Zlotkin (1994)."*

Today in many markets in the financial sector, such as the stock market, already most transactions are conducted by algorithms. In consumer markets, a rapidly increasing share of consumer choices is based on the information provided by AI. In the market of consumers' internet search, AI-based search has in practice a 100 per cent market share. Casual observation of trends suggests that, in one way or the other, market choices by firms and households are increasingly AI-based.

Some 20 years after the above observation, Varian (2016) makes the following observations:

> ***"Personalization and customization:*** *Computer mediation allows services that were previously one-size-fits-all to become personalized to satisfy individual needs. Today we routinely expect that online merchants we have dealt with previously possess relevant information about our purchase history, billing preferences, shipping addresses, and other details. This allows transactions to be optimized for individual needs.*
>
> ***Experimentation and continuous improvement:*** *Online systems can experiment with alternative algorithms in real time, continually improving performance. Google, for example, runs over 10,000 experiments a year dealing with many different aspects of the services it provides, such as ranking and presentation of search results. The experimental infrastructure to run such experiments is also available to the company's advertisers, who can use it to improve their own offerings.*
>
> ***Contractual innovation:*** *Contracts are critical to economic transactions, but without computers it was often difficult or costly to monitor contractual performance. Verifying performance can help alleviate problems with asymmetric information, such as moral hazard and adverse selection, which can interfere with efficient transactions. There is no longer a risk of purchasing a "lemon" car if vehicular monitoring systems can record history of use and vehicle health at minimal cost."*

Parkes and Wellman (2015) suggest that the aims of AI and economics are fundamentally the same. An AI agent's situation aligns squarely with an agent in a standard economic model. According to Parkes and Wellman, AI strives to construct a synthetic *machina economicus*. Research is producing rapidly algorithms that produce efficient resource utilization under rather weak assumptions, see, e.g., Ma et al (2016).

In a dynamic imperfect information game setting, AI algorithms are by now able, see Sandholm (2017), to augment the precomputed blueprint over time to play even closer to Nash equilibrium based on what holes the opponents have been able to identify and exploit. This is in contrast to previous approaches to learning in games, where the goal has typically been opponent modeling and exploitation. Those pervious approaches tend to open the agent up to counter-exploitation and causes the strategy to be opponent specific. In contrast, the self-improvements in the most sophisticated algorithms are universal. Moreover, such algorithms are domain independent. Firms and households can use them to take optimizing decisions in different economic settings.

The basic claim of this paper is that the use of AI will in general improve market efficiency. The claim is based on a simple assumption that without AI, economic agents will take less rational decisions in an imperfect information dynamic setting with huge number of different states of the world-based on the vast amount of rapidly evolving data that is usually available than with AI.

Proposition 1: *AI will, in general, improve market efficiency.*

Proposition 1 holds if market participants are, in general, able to process and learn from available data to take fully rational optimizing decisions only if they use AI. Proposition 1 holds more specifically only if in the Arrow-Debreu model meaning of rationality the following assumptions hold:

i) Without AI market participants are unable to take rational decisions.
ii) With AI market participant are able to take rational decisions.

Regarding assumption i), it is rather apparent that, in general, no market participant has the capacity to process and learn from all the available data to take rational decisions, without resorting to AI.

Regarding assumption ii), Geanakoplos (1987) assesses the meaning of the rationality assumption as a requirement of the efficiency result of the Arrow-Debreu model. In that assessment Geanakoplos (1987, p. 117) rather well specifies that computers should be capable to today's AI, for the rationality assumption of the Arrow-Debreu model to hold, when he notes the following:

> *"As an instance of this last case, note that it follows from the rationality hypothesis that the surge in the microcomputer industry influenced consumer choice between typewriters and word processors only through availability (via the price), and not through any learning effect. (Consumers can 'learn' in the Arrow-Debreu model, e.g. their marginal rates of substitution can depend on the state of nature, but the rate at which they learn is independent of production or consumption – it depends on the exogenous realization of the state. We shall come back to this when we consider information.) If for no other reason, the burden of calculation and attention which rational choice over consumption plans imposes on the individual is so large that one expects rationality to give way to some kind of bounded rationality in some future general equilibrium models."*

With the availability of AI, the deviation from the Arrow-Debreu model rationality has clearly diminished. However, it is beyond the scope of this paper to assess how far today's AI is from Arrow-Debreu rationality.

In addition to the rationality assumption, the efficiency result of the Arrow-Debreu model is based on the symmetry and availability of full information assumptions. However, although the increasing data availability, which market participants can use with the help of AI, should reduce the asymmetry of information, information will always be to some extent asymmetric.

While the more extensive use of data by market participants should also enhance market efficiency, it should be sufficient for the improved market efficiency result that with AI market participants are able to take rational decisions while they are not able to do that without AI.

Proposition 1 holds also only in general because in specific cases AI can also reduce market efficiency. Proposition 3, in section 5, deals with the case where AI aggravates an existing distortion. Moreover, AI can create totally new distortions, as in the case of high frequency trading and the so called arms race on speed examined by Budish et al (2015). This is also discussed in section 5.

4 The effect of AI on labor markets via improved product and capital market efficiency

It is apparent that expansion of digital data and learning from that data by AI makes the market for many products more efficient and competitive. Given this, it follows from the Hicks-Marshall derived factor demand rules that also AI increases labor demand elasticity. The higher the labor demand elasticity, the more the exiting labor market distortions impair an efficient allocation of labor, i.e. the higher is the loss in terms of output and employment.

Increased labor demand elasticity decreases the wage share of income because the incidence of labor taxes shifts from capital to labor and the bargaining power of capital over rents increases.

While AI improves, in general, product and capital market efficiency, it is unlikely to solve the informational problems for labor market entrants. The of lack of job market data on new labor market entrants in combination with more efficient product and capital markets impairs labor market functioning for those persons.

This labor market impairment may arise also from the signal information on long-term unemployed persons (e.g. Kroft et al, 2010) when other markets function more efficiently. The reasons for these impairments are the second best setting in which improved product and capital market efficiency aggravates the effects of existing labor market distortions. The reason for this impairment is the second best setting in which improved product and capital market efficiency aggravates the effects of existing labor market distortions.

Proposition 2: *Improved product market efficiency under AI increases the derived factor demand elasticities. Similarly, improved capital market efficiency increases the elasticity of labor demand. The increased elasticity aggravates the effects of existing labor market externalities such as the scarce information on job market information of labor market entrants.*

Proposition 2 is based on standard elasticity results. AI increases factor demand elasticities through two channels: increased product market efficiency and competition and increased capital market efficiency. Increasing product market competition means that price elasticity of product demand increases. The Marshall-Hicks derived factor demand rule states that: *"The demand for anything is likely to be more elastic, the more elastic is the demand for any further thing which it contributes to produce."* As Kennan (1998) notes, Hicks formulated this rule in *The Theory of Wages*. Kennan (1988) provides the proof that this rule holds.

AI increases also capital market efficiency making capital allocation more fluid or more adjustable. It follows from the Le Chatelier-Samuelson principle that the demand elasticity of a factor (labor) increases when the other factor (capital) becomes more adjustable. Milgrom and Roberts (1988) provide both the proof and a discussion of the LeChatelier-Samuelson principle.

In the case of more efficient capital market, the intuition of the principle can be explained as follows. As capital becomes more adjustable, the derivative of compensated labor demand with respect to its own price is larger than when capital is less adjustable. In other words, with improved reallocation of capital from less to more productive use, labor demand elasticity increases. This is true whether labor and capital are substitutes or complements.

It is a standard result of the theory of the second best that a move from a second best setting with inefficient markets to efficient markets will, in general, aggravate effects of remaining distortions such as the effects of distortionary taxes.

5 The effect of AI on the labor market via skill demand change and increased asymmetry of information

The analysis on high frequency trading by Budish et al (2015) shows that AI, in the case of very fast automated trading in a continuous – as opposed to discrete - time setting, gives rise to a market distortion. In his lecture at the 2017 annual meeting of the American Economic Association, Eric Budish referred to further work suggesting that the distortion would not be, in general, self-correcting.

This analysis focuses on how AI may aggravate distortions via information asymmetry in labor market. This is likely to be the case if AI stresses the importance of non-cognitive skills in relation to cognitive skills which are easier to screen through a potential worker's educational attainment. Recent evidence by Edin et al (2017) and Jokela et al (2017) suggest that the role of non-cognitive skills has already increased.

Proposition 3: *AI aggravates labor markets distortions by increasing asymmetric information when AI increases the role of non-cognitive skills in relation to cognitive skills resulting in a reduced signaling value of educational attainment. This may also increase wage inequality.*

Proposition 3 is based on an assumption that AI increases asymmetric information when AI increases the role of non-cognitive skills in relation to cognitive skills resulting in a reduced signaling value of educational attainment. This assumption, in turn, is based on an assumption that AI is able to perform many of the cognitive tasks that so far only humans have been able to do. This increases the demand for non-cognitive skills in relation to cognitive skills. Employers will learn the true non-cognitive skills of workers only over time.

Proposition 3 is based on standard results in a labor market model with asymmetric information when a firm cannot observe worker's ability at the time of hiring. The firm learns this ability only over time.

Carrilo-Tudela and Kaas (2015) show that when the learning rate is sufficiently low, an equilibrium emerges in which low-wage firms attempt to hire both low- and high-ability workers by offering incentive compatible separating contracts. These contracts offer initially a low wage and then promote the worker by increasing his wage to marginal productivity. High-wage firms offer contracts that pay the same starting wage to all workers, but after learning their employees' types they promote high-ability workers and demote low-ability workers. Low-ability workers can experience a wage cut or a wage rise when undertaking a job-to-job transition, while high-ability workers can only experience wage rises when changing jobs without an intervening spell of unemployment, i.e. wage inequality between high- and low-ability workers increases.

Skill-biased technical change by itself is likely to widen the dispersion of abilities among workers. Hence a new technology, such as AI, is likely to play a role in increased sorting of employees by firms. There is likely to be two also sided heterogeneity (employees' ability and firms' productivity) and two sided asymmetric information (the true ability of the employee and the true productivity of the firm). For both sides there is hidden action and information in line with the analysis in Hart and Holmström (1987).

There is evidence that the increased sorting of workers into high productivity and high wage firms is less than before based on the directly observable ability signals such as educational attainment, see Grossman et al (2015). This stresses the informational problems that may impair the functioning of the labor market.

An obvious general policy conclusion from propositions 2 and 3 is to reduce existing labor market distortions. An equally obvious policy conclusion is to boost skills in the face of changing skill requirements due to AI. These policies would apparently be warranted even in the absence of AI. Based on Finnish and Swedish evidence, non-cognitive skills appear to have become more important already now. Moreover, those policies are likely to be policies that would facilitate sustained job creation in the face of AI's technological effects via automation. The policies are a subject of a separate paper. That paper will also show how unemployment insurance can be reformed so as to alleviate the adverse effects in propositions 2 and 3.

6 Discussion

How realistic are the above-mentioned results? Answer to this question could be given by empirical research. It is possible to evaluate AI's effect on competition in a given market by examining how, e.g., markups and price elasticities change. This econometric work is made more difficult by the fact that it may be impossible to isolate the effect of AI from, say, the effect of increased price transparency due to internet. However, the policy implications are unlikely to differ much whether the source of the increased efficiency and competition is due to AI or due to digitalization including AI.

Based on empirical research on how much price elasticity of a certain product has increased, one can assess how labor demand elasticity in the production of that product has increased. Analytically the Hick-Marshall rule means that the latter elasticity is proportional to the former. Based on the information regarding the change in labor demand elasticity, one can assess how much more, say, a specific tax distortion affects behavior.

In the absence of substantial empirical evidence on AI's effects on, say, consumer markets, one can, however, observe that consumers are already using AI on a large scale. Casual observations support the view that markets for many products are by now fairly competitive. Tax-free prices for a given good do not deviate much. If the AI proposes a choice on a website that offers a much lower price than other sites, one can suspect that the product is a fake one.

AI and internet seem to have increased competition also in the service markets. Search engines are already relatively good at finding, say, films, restaurants, hotels that meet individual preferences at a presumably fairly competitive price. Information is available on, e.g., the quality of individual Uber drivers. Drivers, restaurants, shops and other service providers know that the quality of their services are rated by random customers and that consumer choices are increasingly based on a careful evaluation of all available data.

Given the complexity in arriving at a rational choice that takes into account one's preferences, budget constraint and the vast amount of information on possible products, it is rational for the consumer to take his or her decision using AI. This new division of labor between the computer and the consumer, where consumers are able to take rational decisions with the help of AI, should improve market efficiency.

Firms may be even faster than households to optimize their decisions by using AI. It is evident that, say, an efficient management of global value chains is very difficult without AI. Firms' decisions to locate different parts of R&D, marketing and production chains will take fully into account, in addition to direct costs, the functioning of the labor market, the availability and quality of human capital and infrastructure as well as the stability and quality of governance in different countries.

What could be the implications if all decisions by households and firms would be based on the advice of a rational AI that takes into account all available information? For economic analysis, one implication could be that DSGE models become more relevant than they have been so far.

In practice, more efficient markets with rational firms and households could be expected to greatly aggravate the effects of remaining economic distortions and market failures. The cost of inefficient economic structures could be much higher than today. There could be more crisis countries depending on how fast national policies adapt to the new environment where market decisions are AI-based.

A global economic environment with more efficient markets will be more demanding for policy making. The cost of bad policies will be higher and their effects will presumably be transmitted faster than before. This could shorten the usually long policy lags and lead to more proactive policies.

7 Concluding remarks

AI will, in general, improve efficiency of markets. This holds if, in general, AI is required to reach fully rational decisions based on all available data. The increased market efficiency will aggravate effects of existing economic distortions. AI may also induce new distortions.

The paper puts forward two propositions on how AI affects employment via the market effects. One proposition considers the effects of AI through more efficient product and capital markets under unchanged asymmetric information on labor market entrants. AI is unlikely to solve the informational problems for labor market entrants. The of lack of job market data on new labor market entrants in combination with more efficient product and capital markets impairs labor market functioning for those persons.

This labor market impairment may arise also from the signal information on long-term unemployed persons when other markets function more efficiently. The reasons for these impairments are the second best setting in which improved product and capital market efficiency aggravates the effects of existing labor market distortions. The reason for this impairment is the second best setting in which improved product and capital market efficiency aggravates the effects of existing labor market distortions.

The other proposition considers the effects of AI through changes in the skill requirements from cognitive to more non-cognitive skills making educational attainment a less valuable signal of skills.

The analysis of the paper should be enhanced in many ways. For example, empirical analysis should be undertaken regarding this paper's central hypothesis of AI's market efficiency improving general effect.

There is also a clear need for analysis of potential policies to address the adverse effects of AI. The paper argues that AI impacts the economy, via the markets, earlier than it automates jobs on a larger scale. AI's potential economic effects and their likely speed stress the need for proactive policies.

The paper argues in brief that required policies would amount to addressing existing economic distortions and to boosting skills. These policies would be warranted even in the absence of AI given that there is evidence that the role of non-cognitive skills has already increased. Moreover, the needed policies are likely to be policies that would facilitate sustained job creation in the face of AI's technological effects via automation. A separate paper will analyze policies which would alleviate the adverse labor market effects of AI.

References:

Budish, Eric, Peter Cramton and John Shim (2015), The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Quarterly Journal of Economics* 130(4), Nov 2015, pp 1547-1621 http://faculty.chicagobooth.edu/eric.budish/research/HFT-FrequentBatchAuctions.pdf

Carrilo-Tudela, C. and Kaas, L. (2015), Worker Mobility in a Search Model with Adverse Selection", *Journal of Economic Theory,* Vol. 160 (December) pp. 340 – 386. https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbnxsZW9qa2Fhc3xneDoyZDg4NmQxYmIxODYyM2E

Davis, Steven J. and Haltiwanger, John (2014), Labor Market Fluidity and Economic Performance. NBER Working Paper No. 20479. http://www.nber.org/papers/w20479

Edin, Per-Anders, Peter Fredriksson, Martin Nybom and Björn Öckert (2017), The Rising Return to Non-cognitive Skill. http://www.ne.su.se/polopoly_fs/1.327441.1492512829!/menu/standard/file/Noncognitive_skill_11_PF_MN.pdf

Geanakoplos, John (1987), "The Arrow-Debreu Model of General Equilibrium," in J. Eatwell, M. Milgate, and P. Newman, eds., *The New Palgrave Dictionary of Money and Finance,* Macmillan Press, London, 1992, , pp. 59-68. Also appeared in J. Eatwell, M. Milgate, and P. Newman, eds., *The New Palgrave Dictionary of Economics*, Macmillan Press, London, 1987, Vol. 1, pp. 116-124. Cowles Foundation Paper No. 1090 http://dido.econ.yale.edu/~gean/art/1987-newpalgrave1.pdf

Grossman, Gene, Elhanan Helpman, and Philipp Kircher (2015), "Matching, Sorting, and the Distributional Effects of International Trade." (forthcoming) Journal of Political Economy. http://scholar.harvard.edu/files/helpman/files/matching_and_sorting.pdf

Hart, Oliver and Bengt Holmström (1987), The theory of contracts. In Bewley, T. Advances in Economics and Econometrics. Cambridge University Press. pp. 71–155. https://dspace.mit.edu/bitstream/handle/1721.1/64265/theoryofcontract00hart.pdf?sequence=1

Jokela, Markus Tuomas Pekkarinen, Matti Sarvimäki, Marko Terviö ja Roope Uusitalo (2017), Secular rise in economically valuable personality traits. PNAS, vol. 114 no. 25, 6527–6532. http://www.pnas.org/content/114/25/6527.full

Kennan, John (1998), The Hicks-Marshall Rules of Derived Demand: An Expository Note. http://www.ssc.wisc.edu/~jkennan/teaching/hicks_marshall.pdf

Kroft, Kory, Fabian Lange and Matthew J. Notowidigdo (2013), Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment. *Quarterly Journal of Economics* (2013) 128 (3): 1123-1167. https://academic.oup.com/qje/article-lookup/doi/10.1093/qje/qjt015

Ma, Hongyao, Reshef Meir, David C. Parkes and James Zou (2016), Contingent Payment Mechanisms to Maximize Resource Utilization. http://econcs.seas.harvard.edu/files/econcs/files/ma_16contingent.pdf

Milgrom, Paul R. and John Roberts (1996), The LeChatelier Principle. *American Economic Review*. March 1996, Vol. 86, Issue 1, Pages 173-179. http:/web.stanford.edu/~milgrom/publishedarticles/The%20LeChatelier%20Principle.pdf

Parkes, David C. and Michael P. Wellman (2015), Economic reasoning and artificial intelligence. *Science* 17 Jul 2015: Vol. 349, Issue 6245, pp. 267-272.
http://science.sciencemag.org/content/349/6245/267/tab-pdf

Sandholm, T. (2017), Super-Human AI for Strategic Reasoning: Beating Top Pros in Heads-Up No-Limit Texas Hold'em. Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17).
https://www.ijcai.org/proceedings/2017/0004.pdf

Stanford University (2016), One Hundred Year Study on Artificial Intelligence (AI100). 2016 Report.
https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf

US Government (2016), Preparing for the Futue of Artficial Intelligence. Executive Office of the President National Science and Technology Council Committee on Technology.
https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

Varian, Hal R. (1995), Economic Mechanism Design for Computerized Agents. WOEC'95 Proceedings on the 1st conference on USENIX Workshop on Electronic Commerce – Volume 1.
http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.56.1930&rep=rep1&type=pdf

Varian, Hal R. (2016), Intelligent Technology. Finance & Development, September 2016, Vol. 53, No. 3, IMF.
http://www.imf.org/external/pubs/ft/fandd/2016/09/varian.htm

Walsh, Toby (2016), The Singularity May Never Be Near.
https://arxiv.org/pdf/1602.06462.pdf

Walsh, Toby (2017a), Expert and Non-Expert Opinion about Technological Unemployment.
https://arxiv.org/ftp/arxiv/papers/1706/1706.06906.pdf

Walsh, Toby (2017b), The AI Revolution. An essay commissioned by the NSW Department of Education.
https://education.nsw.gov.au/media/exar/The_AI_Revolution_TobyWalsh.pdf