Economic Decision Making, Present and Future

Charles Noussair[1]

1. Introduction

Economic decision making requires intelligence. Intelligence is the ability to learn, reason, solve problems, perceive, comprehend language, and adapt behavior to fit new circumstances. One can place all intelligence into one of the two categories. The first is biological intelligence, possessed in some measure by all animals, and which is found in nature. The second, artificial intelligence, is man-made. The capability of biological intelligence is more or less fixed, but that of artificial intelligence is growing exponentially over time.

Several decades ago, all important economic decisions were made by biologically intelligent beings. Over time, more and more decisions are made by artificial agents such as algorithmic traders and control systems. In the not-to-distant future, artificial decision makers, in the form of computer programs, robots, or artificially enhanced human beings, will dominate the economy.

For the moment, technological improvements are necessary to expand the capabilities of artificial intelligence. That is, humans must make scientific progress for artificial intelligence to become more powerful. However, the future holds the possibility of self-enhancing artificial intelligence, which promises to improve these capabilities on its own. What we can expect over time is clear. An increasing percentage of economic decisions will be taken by artificial entities, and their capabilities will soon vastly outclass those that biological entities have. Economic science, like society as a whole, must adapt to this new stage of history.

2. Biological Intelligence

As a behavioral scientist, I have spent the last three decades studying human decision making, that is, biological intelligence. This is what is investigated by behavioral and neuro scientists. We are gaining a good understanding of how biological intelligence operates. The main message from behavioral experiments for economists is that individuals' decisions depart from what economists typically consider rational behavior. Thus, biological intelligence, while reasonably good at making many types of decisions, is characterized by a degree of inconsistency and suboptimal decision making.

---

[1] Department of Economics, and Director, Economic Science Laboratory, University of Arizona. Correspondence to cnoussair@email.arizona.edu. These remarks were delivered as the 2017 Mabel F. Timlin Lecture at the University of Saskatchewan on March 8, 2017. The author thanks the Timlin Trust for making the lecture possible. I am greatly indebted to my colleague Ad van de Gevel, who initially sparked my interest in the connection between artificial intelligence and economic behavior. Many of the ideas expressed in this lecture are due to him. The material in this lecture is based on the book "The Nexus between Artificial Intelligence and Economics" by Ad van de Gevel and Charles Noussair (2013).

Many departures from rationality can be classified into several prominent patterns. Some decisions are simply too difficult because they require a level of computation that it either beyond the capability or prohibitively costly for a human to perform. For example, the mathematical problem, "What is the cube root of 7,750,636,739?" is beyond the ability to most humans to solve. If confronted with this problem, some of us would be totally lost. Some of us would try to approximate it. Still others of us would become exhausted trying to compute it.

Some decisions are too tricky and trip up most people. Consider for example the Cognitive Reflection Test questions analysed by Frederick (2005). An example of one of these questions is the following "A patch of lily pads in a lake doubles in size every day. If it takes 48 days for the patch to cover the entire lake, how long does it take it to cover one half of the lake?" At first glance, it would seem that the answer would be 24, and may individuals answer that way. Their rationale is that it should take half of the time to cover half of the area. Upon reflection, however, one should realize that the correct response is 47, since the lily pad's size will double from day 47 to day 48. Questions of this type are very tricky since the first response that comes to mind often seems obviously correct to the individual, but in fact it is wrong.

Some decisions require more patience or willpower than the typical person has. In Tucson, Arizona, a standard price for a payday loan is $15 - $40 for a loan of $100 for two weeks. The implied annual interest rate for the low fee of $15 is 1,433%. Often, a lack of patience or willpower gets one into a position where they have to take out such a loan.

Sometimes emotions influence decisions, and in some cases this can be harmful to an individual. Let me describe two examples from my research of how emotions can adversely affect outcomes. Breaban and Noussair (2018) track the emotional state of traders as they participate in experimental asset markets. In this experiment, subjects are traders who can buy or sell units of an artificial, dividend-paying asset. Subjects are compensated financially based on their profits in the experiment, and so they have incentives to trade profitably. The authors found that those individuals who kept their emotions in check, particularly at times of market volatility, made more money.

Dalton et al. (2017) studied the link between emotions, poverty, and productivity. They conduct an experiment where they show subjects a video, and then ask them to complete a task. The task involves setting the highest possible number of sliders, which are displayed on the subjects' computer screen, in the exact middle point of a pre-specified range, using their computer mouse to move a cursor (Gill and Prowse, 2012). There are two treatments, the *Poverty Treatment*, in which participants watch a video of a group of poor people who are living in a garbage dump and scavenging for food, and another control condition called the *Neutral Treatment*, in which subjects watch an informative video about a national park in Alaska. Participants who viewed the poverty video in the Poverty treatment did not perform as well on the task as the control group in the Neutral treatment. The emotions invoked by the poverty video accounted for some of the difference.

Classical economics assumes that we have none of these problems. It supposes that we can make computations of unlimited complexity, recognize and adapt to tricky situations, have as much patience and willpower as we need, and are not influenced by emotions. In contrast, behavioral economics recognizes these limitations and tries to analyze them.

Neuroscientists have uncovered some of the physiological processes underlying biological intelligence. The brain has many functional neural pathways, which transmit electrical signals to, from, and within the brain. These pathways consist of neurons, along which the signals travel. Along the pathways, there are gaps between the neurons, and these gaps are called synapses. The synapses contain chemicals called neurotransmitters, which are released into the synapses to transport electrical signals to the downstream neurons and are then reabsorbed. A few of the pathways consist of dopamine neurons, which are characterized by their ability to release dopamine into the synapses along the pathways. Dopamine neurons fire when you do something pleasurable. Your brain then seeks ways of getting them to fire again. It engages areas involved in planning and social cognition to do so. In this manner, you learn what is valuable or brings pleasure to your organism. This is a very inefficient process of learning compared to what exists in artificial forms today.

How efficient is the human brain in computation? Biological human thinking is limited to 1016 calculations per second per human brain. All thought is taking place on neurons with a top speed of electrochemical signals of 150 meters per second along the fastest neurons. This peak is fleeting; we lose neurons as we age and their capacity declines over time. In comparison, at the time of this writing, speeds in modern computer chips are currently at around 2GHz, a ten million fold difference over humans. The difference is increasing exponentially.

3. Artificial Intelligence and its impact

A look at world GDP from 5000 BC until the present reveals that the industrial revolution was a singular event in world history, as the world economy shifted from agriculture to industry. It changed the doubling time of world output from every 900 years previously to every 15 years subsequently. Could a new transition be imminent?

Singularitans, the most prominent of whom is Ray Kurzweil, claim that a similar transition will occur at some time between 2045 – 2140. In their view, a period of infinite progress (or at least progress that is much more rapid than has ever occurred previously) would set it after that date.

What might cause such a singularity? It would have to be something as sweeping as industrialization, which would affect the entire economy. There are a number of industries in which technical progress is proceeding rapidly, such as biotechnology, telecommunications, and nanotechnology. However, advances in any of these areas increase productivity in only a few sectors. The only potential candidate, in my view, for providing a paradigm shift on the order of the industrial revolution, is self-enhancing artificial intelligence. Self-enhancing refers to the ability to improve on its own. Artificial intelligence with this capability could drastically affect the productivity of the entire labor force and capital stock.

Artificial intelligence is of course already here and is more and more of a factor in our economy. The artificial decision makers in our economy can take three forms. The first are computer programs, ubiquitous in the modern economy. We can expect these computer programs to become more intelligent in the future. The second is embodied artificial intelligence, in the form of robots, and the type and number of these will also only increase. The third are augmented humans, people who are implanted with artificial devices that enhance their mental capabilities. This is on its way and will likely take two forms initially. One is the implantation of nanobots into the human brain, and the other is improvement in Brain Computer Interface (BCI).

Nanobots will augment our brain with nonbiological intelligence, starting with routine functions of sensory processing and memory, and then moving on to skill formation, pattern recognition, and logical analysis. The technology is progressing rapidly. BCI refers to a direct connection between a human (or animal) brain and an external device. These connections range from non-invasive technologies that recognize brain signals externally, to invasive technologies that involve surgery and direct electrode implantation.

Whole Brain Emulation (WBE) is the transfer of a mind, the mental structure and consciousness of a person, from a biological brain to an external carrier, such as a computer. The word emulation describes the attempt to achieve as close a functional match as possible to the source biological brain, so that the mind is altered as little as possible in the transfer. It is recognized that the copied mind would not be identical to the original, since so much of human thinking is directed towards its physical needs and environment, and these depend on characteristics of the host body. A successful emulation need not predict all details of the original behavior of the emulated system; it need only replicate computationally relevant functionality at the desired level of emulation.

There are specific technologies that are enabling the advancement of artificial intelligence. In particular, several precursor technologies are currently undergoing revolutions. The first is genetics, and in particular, regenerative medicine and enhancement. The second is nanotechnology, which is enabling the redesign and building of brain and bodies. The third is robotics, which is advancing rapidly and is crucial to the design of embodied artificial intelligence.

4. Implications of Artificial Intelligence

The advent of artificial intelligence portends a fundamental transformation of our society, and indeed of our species. We can think of the next stage of our civilization as one of Robo-Sapiens, in which there is no clear distinction between human and machine. Indeed, we would have a human-machine civilization, characterized by brain computer interface, nanobots in the brain, artificial organs and body parts, and whole brain emulation.

Philosophers will have to revisit the question of what it means to have consciousness. There is no consensus among scientists about the nature of consciousness of non-human entities. Our future non-biological replicas will be vastly more intelligent than we are and therefore could exhibit the finer qualities of human thought to a very large degree. These non-biological entities could likely convince other humans that they are conscious using the emotional cues that humans employ, so that, regardless of whether or not they actually have self-awareness, they will be indistinguishable from conscious beings.

The possibility of uploading one's consciousness onto a permanent platform which can be backed up, evokes the theological concept of the soul, which lives beyond the death of the body that hosted it. Debate about how to interpret this form of eternal life in the context of religious teaching may require religious authorities to offer pronouncements on the matter.

Artificial intelligence offers the possibility of very long increases in lifespan, which De Grey (2008) has termed Methuselarity. Currently, human life expectancy is increasing by .2 - .3 years every year on average in the developed world. If this factor reaches 1, longevity escape velocity (LEV) will be reached. In another words, the longer an individual has lived, the longer she can expect to continue living into the future. De Grey argues that only modest advances in medical technology are required to achieve LEV. Some of the medical and diagnostic innovations that would underpin this progress are smart toilets, databases with individual DNA, organ growth and replacement technologies, and regenerative medicine. De Grey asserts that the first 1000 year old will be born less than 20 years after the first 150 year old. In the not-to-distant future, most deaths will be caused by accidents or catastrophic events rather than natural causes.

Despite its promise, there are serious obstacles to realizing these new advances in artificial intelligence. First, as Jones (2005) has argued, there is a ”Knowledge Burden”, or information overload problem. Scientists will need ever more training to reach the research frontier as more knowledge is accumulated by humanity. This burden acts as a brake on the acceleration of technological and economic progress. The available evidence is that the knowledge burden has been increasing over time, with negative consequences for economic growth.

Political considerations may get in the way of scientific advances. Tension between those who benefit from new technologies and those who are hurt may result in policies that slow down innovation. Beneficiaries will seek to avoid compensating those who are hurt, while the latter will demand compensation for their losses or a reinforcement of the status quo. In addition, psychological factors might delay or preclude the onset of a singularity. Human resistance to, and fear of, change can be powerful. It may also the case that, when basic human needs are already satisfied, the motivation to innovate may decrease. There are also issues of coordination, as multiple technologies must all advance for a singularity to be reached. There are also economic incentives to slow down innovation on the part of the providers and owners of incumbent technologies. Producers of goods and technology tend to resist making their own products obsolete.

Nonetheless, let us suppose for the purpose of this lecture that these obstacles are not overwhelming and let us imagine a future in which innovations in artificial intelligence advance rapidly and diffuse thoroughly into the economy.

Numerous new legal and ethical issues will arise. I believe that legal issues revolving around the rights and responsibilities of artificial agents will be some of the most important and contentious debates among legal scholars this century. What rights will intelligent robots be allowed to have? Should they be able to purchase their freedom from their owners? Will robots be allowed to own other robots or other types of property? If a robot becomes indistinguishable from a biological human, should it be accorded the right to vote? Should it be required to pay taxes on its income?

Should there be laws that protect robots from being required by their owners to perform dangerous tasks that might damage them? Do owners have a legal obligation to act as a protective guardian for the robots they own as they do for their own biological children?

There is a risk that artificial intelligence will turn on its human creators. Bostrom (2002) writes that artificial intelligence will have the capability to bring about human extinction. Even if the intelligence is not hostile, its potential indifference to humans could make it very dangerous. Furthermore, the military applications of artificial intelligence can potentially lead to intentional apocalyptic violence and destruction. A basic question arises regarding the appropriate laws of war for this new age: will a robot be allowed to make the decision on its own to kill opposing forces? Artificial intelligence has already become a dimension of great power competition. As noted by Russian President Vladimir Putin in 2017, "Artificial intelligence is the future, not only for Russia, but for all mankind. … Whoever becomes the leader in this sphere will become the ruler of the world".

As a precaution, some have argued that friendly artificial intelligence must be developed first. Eliezer Yudkowsky (2008) of the Singularity Institute for Artificial Intelligence has called for the creation of "Friendly AI", smart enough to improve on its own source code without programmer intervention, to mitigate the existential threat of hostile intelligences. The primary objective of this friendly artificial intelligence would be to protect humans. A starting point for the code of conduct for friendly artificial intelligence is Isaac Asimov's *Three Laws of Bobotics*. These are:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

While the field of artificial intelligence is not advanced enough to be sure that Friendly AI can be created, there are strong arguments for doing so if it turns out to be feasible.

5. Economic Analysis

What can economic theory tell us about how artificial intelligence will affect our economy and society? Let me offer a few observations.

Robots will certainly replace humans in many professions and this process has already begun. To identify the jobs that robots will do and which will be left to humans, a basic principle called the Law of Comparative Advantage can be applied. Essentially, we can place every job on a spectrum that can be ranked by how well humans (possibly augmented ones) perform compared to robots. The Law says that humans will continue to do those activities in which they are relatively less inefficient. This means that no matter how productive robots become, human labor would still be used in jobs in which the robot advantage is smaller than for other jobs. The consequence is that robots cannot be expected to knock humans out of every job. While this may be surprising to some, it is supported by the sweep of human history. At no time in the past has a labor-saving technology ever led to permanent unemployment.

Where different jobs fall on the spectrum may be surprising. Robots can obviously vastly outperform humans in repetitive physical tasks such as assembly line work. However, Moravec's (1988) paradox reveals another potent pattern: this is that high level reasoning is relatively easy for computers compared to humans, while computers have relatively weak sensorimotor skills. Thus, tasks that involve high level reasoning will be among the tasks from which humans are most readily displaced. As Pinker (2007) writes "… it will be the stock market analysts and petrochemical engineers that are in danger of being replaced by machines. The gardeners, receptionists and cooks are secure in their jobs for decades to come". Humans have a comparative advantage in these non-repetitive physical or creative tasks.

How much will humans earn for their work and how will the prices change for the goods that they purchase? Prices for most consumables will be much lower than they are now because of greater efficiency in production. This means that the average person will be able to consume more.

Human wages will be dependent on robot wages, since they will both participate in the same labor market. Robot wages can be determined with economic theory using the Iron Law of Wages (Lassalle, 1863). Lasalle's 19th century insight was that if new labor can always be supplied, wages will eventually equal the cost of subsistence of the labor. The idea was originally formulated in response to the industrial revolution, when it was thought that there was essentially an infinite number of surplus rural workers who could move to the cities and work in the factories. Competition between workers for a finite number of jobs would ensure that wages were bid down to the minimum that it would take to get the worker to work, which is the minimum subsistence level required to keep that person in good enough health to perform their job. As it turned out, the supply of workers that appeared to be infinite at the outset of industrialization was eventually exhausted, and wages did begin to rise above subsistence.

The prospect of the mass production of robots may give new life to Lasalle's theory. If industrial robots can be mass produced at constant marginal cost, their potential supply would be essentially infinite, or at least enough to satisfy demand for them at the marginal cost of production. The wages to robots, in the sense of the cost to the firm to employ them in production, would equal the maintenance cost of keeping the robots running.

How human wages would change would vary by individual and depend on whether the human in question has become more or less valuable in a world with robot labor. The economic value of neural augmentation may be considerable, and people with such augmentation may command higher incomes. Individuals who support the new industries, such those who design, program, build and maintain robots or non-embodied types of artificial intelligence, may become very valuable. On the other hand, those humans that are in industries in which robots have a comparative advantage are likely to suffer a decrease in their income. This decrease may or may not be fully offset by the lower prices of consumer goods, and such individuals may or may not experience an improvement in their standard of living.

The fate of human workers also hinges on how strong preferences are for human labor. Wealthy people may prefer to have natural humans rather than robots as their household employees and caregivers. Some clients may prefer dealing with human sales people and service personnel to interacting with robots. Initial indications are that there is high demand for robots that express emotions, presumably because doing so renders them more like humans.

What will be the fate of academics like me? My job consists primarily of teaching and conducting research. For the moment, academics are not facing competition from artificial agents. However, we are probably less than two decades away from having an artificially intelligent teaching program that can teach economics better than I can, at least at the introductory level. The wait would likely be longer for upper level and graduate courses because the smaller number of students at higher levels of study renders the profitability of developing artificial instructors lower. Small, research-oriented graduate courses will probably be taught by human instructors into the foreseeable future. Whether I will still be teaching other courses will depend on demand for human instructors. If students are willing to pay a premium to have contact with a human, a segment of the market for large introductory courses taught by humans can persist. My refuge will be research, where creativity and original thinking is required. Artificial intelligence may help in generating ideas, identifying related research, data analysis, and solving difficult mathematical problems. However, AI that has all of these capabilities is likely still some decades away. My problem will be competition from humans that are younger and smarter than I.

Methuselarity, large increases in lifespan with relatively short morbidity, would also transform the economic landscape. The percentage of adults who are working should increase since a greater fraction of one's adult life would be healthy. This, along with robot labor, can potentially lead to strong improvements in the dependency ratio. Historically, increases in life expectancy have greatly increased income and there is no reason to suppose that the same pattern would not hold in the future. With very long lifespans, individuals may have multiple sequential careers, with some people retiring from one career and returning to begin a new one, perhaps multiple times and possibly after long interludes between careers.

Some interesting implications may be in store for economic theory. In classical models, agents are assumed to have no computational limits in making the best decisions to achieve their objectives. This ability will be a defining characteristic of intelligent artificial agents, and will switch from being a normative to a descriptive assumption, as artificially intelligent agents make more and more of the economy's important decisions. Classical economic agents are also not taken in by tricky decision problems, have as much patience as needed, and are not swayed by emotions. Whether future AI has these features or not depends on how they are programmed, and I would expect some heterogeneity in this regard. For examples, robot caregivers may be programmed with emotions, while robot back tellers may not be.

In classical dynamic models of savings and economic growth, such as those by Ramsey (1929), Cass (1965) and Koopmans (1963), the economy is modeled as an individual who is assumed to be infinitely lived. Methuselarity could make this assumption more reasonable, as individuals would optimally be planning for very long futures. Furthermore, established models of endogenous growth, such as those of Lucas (1988) and Romer (1986), have as accelerating growth as a fundamental feature. These models can potentially serve as a basis for a line of research describing how a singularity might be achieved and some of its consequences.

It is clear that artificial intelligence will boost economic output. But will it increase happiness? There are at least two ways that an improvement in subjective well-being could occur. One is by increasing incomes. Here, history provides a consistent guide. Past improvement in income over time has not improved average happiness. Indeed, if income inequality increases as well, a boost in average income may actually reduce average happiness, since there is some evidence that income inequality is associated with lower average subjective well-being.

A second possibility is that new technologies will be created as a result of developments in AI that would directly increase happiness. The historical evidence regarding this possibility is mixed. Artificial happiness, created pharmacologically, or cosmetically, has had mixed success. Drugs currently available to improve mood have only temporary success in altering emotional state. Cosmetic surgery does seem to be effective in increasing subjective well-being according to existing survey evidence.

6. Conclusion

Artificial intelligence is increasing rapidly in capability. While there was a time when all economic decisions were made by biological agents, artificial decision makers are becoming more and more prominent. This will have fundamental implications for our economy and also for the science of economics. The economic will be transformed in fundamental ways, but existing economic theory can still be applied to make predictions about our future. Behavioral economics will increasingly have to focus on artificial as well as human decision makers, as well as both types of agent in interaction with one another.

References

Bostrom, N. (2002), "Existential Risks". *Journal of Evolution and Technology* 9, March issue.

Breaban A. and C. N. Noussair (2018), “Emotional State and Market Behavior”, *Review of Finance*, forthcoming.

Cass, David (1965), "Optimum Growth in an Aggregative Model of Capital Accumulation", *Review of Economic Studies* 32 (3): 233–240.

Van de Gevel, A, and C. N. Noussair (2013), “The Nexus between Artificial Intelligence and Economics”, Springer Briefs, Springer Verlag, Heidelberg, Germany.

De Grey, A. (2008), The Singularity and the Methuselarity: Similarities and Differences, in Strategy for the Future, R. J. Bushko ed., IOS press, Amsterdam, The Netherlands.

Dalton, P. S., V. Gonzalez, and C. N. Noussair (2017), “Exposure to Poverty and Productivity”, *Public Library of Science (PLOS-1)* 12(1), 1 - 19.

Frederick, S., (2005), "Cognitive Reflection and Decision Making" *Journal of Economic Perspectives* 19, 25–42.

Gill, D. and V. Prowse (2012), "A Structural Analysis of Disappointment Aversion in a Real Effort Competition", *American Economic Review*, 102, 469–503.

Jones, B. F. (2005), "The Burden of Knowledge and the Death of the Renaissance Man: Is Innovation Getting Harder?" NBER Working Paper 11360.

Koopmans, T. (1963), "On the Concept of Optimal Economic Growth" Cowles Foundation Discussion paper 163.

Lasalle F. (1863), Offenes Antwortschreiben, Leipzig, Germany.

Lucas, R. (1988), "On the Mechanics of Economic Development". *Journal of Monetary Economics* 22, 3–42.

Moravec, H. (1988), Mind Children, Harvard University Press, Boston, MA, USA

Pinker, S. (2007), The Language Instinct, Harper Perennial Modern Classics, New York, NY, USA.

Ramsey, F. P. (1928). "A Mathematical Theory of Saving", *Economic Journal* 38 (152): 543–559.

Romer P. (1986), Increasing Returns and Long Run Growth, *Journal of Political Economy* 94, 1002–1038.

Yudkowsky, E. (2008), "Artificial Intelligence as a Positive and Negative Factor in Global Risk," in Nick Bostrom and Milan Cirkovic, eds., Global Catastrophic Risks, Oxford University Press, Oxford, UK.