

NBER WORKING PAPER SERIES

AI AS THE NEXT GPT:  
A POLITICAL-ECONOMY PERSPECTIVE

Manuel Trajtenberg

Working Paper 24245  
<http://www.nber.org/papers/w24245>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
January 2018

The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Manuel Trajtenberg. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

AI as the next GPT: a Political-Economy Perspective  
Manuel Trajtenberg  
NBER Working Paper No. 24245  
January 2018  
JEL No. J24,O3,O33,O38

### **ABSTRACT**

History suggests that dismal prophecies regarding the impact of great technological advances rarely come to pass. Yet, as many occupations will indeed vanish with the advent of AI as the new General Purpose Technology (GPT), we should search for ways to ameliorate the detrimental effects of AI, and enhance its positive ones, particularly in: (1) education and skills development: revamp the centuries-old “factory model” of education, and develop instead skills relevant for an AI-based economy – analytical, creative, interpersonal, and emotional. (2) The professionalization of personal care occupations, particularly in healthcare and education; these are to provide the bulk of future employment growth, yet as performed today involve little training and technology, and confer low wages. New, higher standards and academic requirements could be set for these occupations, which would enable AI to benefit both providers and users. (3) Affect the direction of technical advance – we distinguish between “human-enhancing innovations” (HEI), that magnify and enhance sensory, motoric, and other such human capabilities, and “human-replacing innovations” (HRI), which replace human intervention, and often leave for humans mostly “dumb” jobs. AI-based HEI’s have the potential to unleash a new wave of creativity and productivity, particularly in services, whereas HRI’s might just decrease employment and give rise to unworthy jobs.

Manuel Trajtenberg  
Eitan Berglas School of Economics  
Tel-Aviv University  
Tel-Aviv 69978  
ISRAEL  
and NBER  
manuel@post.tau.ac.il

## 1. Introduction

AI and related technologies are being heralded as “the next big thing”, one that promises to revolutionize many areas of economic activity and thus to have a profound impact on economic growth. However, the rise of AI coincides with a recent wave of pessimism in terms of productivity growth, expressed forcefully by prominent economists such as Larry Summers (2016) and more thoroughly by Bob Gordon (2016).

Side by side with the gloom, the new “technology enthusiasts” envision a not-too-distant future in which AI will displace most (*all?*) human occupations while unleashing tremendous gains in productivity. This view poses once again disturbing questions about the future of employment, the distributional consequences of mass displacement, and so forth.

Nobody holds the crystal ball, hence rather than arguing about the inscrutable future, it is at least as important to inquiry into what we can learn from history regarding episodes like this, that is, the appearance of a major new technology that is posed to have profound economic implications. Of course, the future is never a replay of the past, but it may provide a useful benchmark against which to assess the unfolding of the new technology.

Mokyr (2017) sounds a cautionary note in that regard: ever since the dawn of the Industrial Revolution in the late 18<sup>th</sup> century, both the pessimists and the enthusiasts have almost invariably proven wrong. Moreover, Mokyr dismantles with solid historical and present day evidence Gordon’s claim that technological advance is bound to slow down in a deterministic fashion (in particular the claim that “all low hanging scientific and technological fruit has already been picked...”).

However, nothing can be taken for granted - as Mokyr skillfully describes, institutions (including government policies) may play a key role enabling or retarding innovation. This is precisely the focus here: given that AI is poised to emerge as a powerful technological force, I discuss ways to mitigate the almost unavoidable ensuing disruption, and enhance AI’s vast benign potential. This is particularly important in present times, in view of political-economic considerations that were mostly absent in previous historical episodes associated with the arrival of new GPTs.

## 2. Is this time different? The political-economy of technological disruptions

The presumption here, well argued in other papers in this conference,<sup>3</sup> is that AI has the potential of becoming a “General Purpose Technology” (GPT) in the foreseeable future,<sup>4</sup> thus bringing about a wave of complementary innovations in a wide and ever expanding range of applications sectors. Such sweeping transformative processes always result in widespread economic disruption, with concomitant winners and losers.

The “winners” are primarily those associated with the emerging GPT sector itself, and those that are at the forefront of the deployment of the GTP in the main applications sectors. They tend to be young, entrepreneurial, and equipped both with the technical knowledge and the

---

<sup>3</sup> See Cockburn, Henderson, and Stern (2017).

<sup>4</sup> See Bresnahan and Trajtenberg (1995).

skills that are made relevant by the new GPT. The labor force composition of Silicon Valley offers a grand view of who are the winners in the present ICT/Internet era. There are further winners in those sectors that are ancillary to the core GTP circle, be it in services that directly benefit from the growth of the GPT (e.g. the VC industry, patent lawyers, designers, etc.), or in others that just ride on the localized boom (e.g. upscale restaurants and entertainment, gyms, tourism, etc.).

The “losers” are mostly those employed in sectors that structurally cannot benefit from the unfolding GPT (“laggards”), and those in industries where the adoption of the new GPT renders many existing competencies and skills obsolete, thus bringing about massive layoffs. They tend to be middle-aged, have lower than average educational levels, and reside in areas that don’t have much diversified sources of employment.

As economists, we tend to view the big sweep of economic growth since the Industrial Revolution as the very embodiment of the “Idea of Progress” (as conceived in the Enlightenment), and hence the rate of growth of GDP as an unequivocal uptick in the welfare of society as a whole. Sure, we do acknowledge that there are distributional consequences, and sure, ever since Pareto we know that we are not allowed to “sum-up utilities” (and thus the “minuses” of losers do not cancel out with the “pluses” of winners). But those half-hearted qualifications become just lip service – the truth is that we rarely dwell into the balance of winners and losers, and in particular we don’t pay much attention to the later. Paraphrasing the well-known dictum of Isaac Newton, we may say that,

*“We enjoy today higher standards of living because we are standing on the broken backs of those that paved the way for technological progress, but did not live long enough to benefit from it.”*

Partly in response to these inequities, the post-WWII era saw the creation of the welfare state, including unemployment insurance, transfers to the disadvantaged, some form of health insurance, re-training programs, etc. These “safety-nets” were supposed to provide a reasonable palliative to “losers”, but the truth is that we still don’t have effective mechanisms to prevent or ameliorate the costs of major technologically-induced transformations.<sup>5</sup> Moreover, existing safety-nets will quite likely fail to cope with the juxtaposition of two new and powerful phenomena: (1) much larger flows of GPT-displaced workers *and* (2) a new “great demographic transition”. Let us examine each in turn.

Regarding the extent of displacement: technological change always causes disruption, as brilliantly articulated by Schumpeter’s notion of “Creative Destruction”. Furthermore, there are inflection points as a new GPT starts working its way through the economy, when in relative short notice very many sectors, competencies and skills became laggards and obsolete.

However, as clearly envisioned in this conference, AI in its various incarnations seems to go much further, in that it has the potential to replace a very wide swatch of human occupations. Many argued forcefully that there are *no* occupations that cannot be eventually

---

<sup>5</sup>Typically, these safety nets function reasonably well when dealing with the consequences of not-too-pronounced business cycles, or with small, *temporarily* deprived groups of the population. Not so when there are major structural transformations, or when the underlying conditions that led to welfare dependency become permanent.

replaced by AI, and that the *vast majority* of present occupations will indeed vanish within a generation.

The consensual view seems to be that a large proportion of employment as we know it today will give way to smart machines, and therefore that  $x\%$  of workers will be displaced, whereby  $x$  is thought to be *significantly larger than in previous GPTs*. At the same time, the extent to which new, presently unforeseeable occupations may arise (denote them  $y\%$ ) seems to be constrained by the very nature of AI: presumably AI will be able to perform also most of the new tasks, and hence they will not constitute a good enough counterbalance to the disappearing jobs, as has been the case in the past. The prevailing view is then that the net displacement of employment ( $x-y$ ) will turn out to be significantly larger for AI than in previous episodes of technological disruption, posing a serious challenge to traditional economic policies.

The second part of the challenge entails a steep drop in birth rates together with the extension of life expectancy (which has been steadily growing for well over a century). These powerful demographic forces have resulted in aging populations, with the concomitant increase in the dependency ratio and the looming threat on the long-term viability of pension system. Notice that life expectancy is now increasing well past the retirement age, so that a typical person in her 50's contemplates a further stretch of 25-30 years of life. Thus, the prospect of being permanently laid off at that stage in life has dire consequences for the displaced individual as much as for society as a whole.

The joint effect of a large influx of displaced workers at the seemingly unique inflection point posed by AI, together with their longer life expectancy, may thus create a formidable challenge to even the most advanced welfare state will be hard pressed to cope with. Put it differently, we cannot afford to have many more, and longer lived unemployed or underemployed people. This is what is at stake with the advent of AI.

There is yet another significant development that magnifies the challenge, and that is the "*democratization of expectations*". The growth in income per capita involves not only a rise in material standards, but in other no less important dimensions of wellbeing, including reduced uncertainty and a concomitant heightened sense of control over our own lives, which entails also the expectation of having a voice in processes that affect us (Hirschman, 1970). Not by coincidence, economic growth and expanding democracy have more often than not gone hand in hand within, as well as across countries.

The Luddites of the early 19<sup>th</sup> century surely had their voice heard, as did their like-minded emulators over the following decades. However, they could hardly expect to make a dent on their fate: democracy was still highly limited and living standards still very low for the vast majority, so that most people were just consumed by the need to provide for their basic needs.

Much has changed since, and nowadays virtually every individual in advanced western countries has come to expect to be entitled, at least in principle, to full participation in every realm of society: the political, the economic, the cultural. The expectation is not just to vote in periodic elections but to have an influence via "participatory democracy"; not just to hold a job, but to partake in the benefits of economic growth – this is what constitutes "the democratization of expectations".

We claim that in such context it has become much harder to have some (many?) bear the costs of technological disruption (the “losers”), while others reap the benefits (the “winners”). Moreover, the “losers” have become much more skeptical of the vague promise that *eventually* the benefits will “trickle-down” to them as well. With good reason: experience shows that the losers typically remain on the downside, even if the welfare state somehow softens their human costs. In advanced, democratic societies, people have become more impatient, more demanding of government, more intolerant of false promises as well as of collective failures. Again, this should be surely considered a highly *positive* by-product of the rise in living standards.

The sharp split between winners and losers, if left to its own, may have serious consequences far beyond the costs for the individuals involved: when it coincides with the political divide, it may threaten the very fabric of democracy, as we have seen recently both in America and in Europe. Thus, if AI bursts into the scene and triggers mass displacement of workers, and demography plays out its fateful hand, the economy will be faced with a formidable dual challenge, that may require a serious reassessment of policy options:

- Governments may have to assume a wider responsibility for navigating effective transitions from old to new GPT’s, and not just for alleviating some of the costs. As said above, the “democratization of expectations” will not allow just for cosmetic adjustments - the political economy of it will eventually force real change.
- In so doing, governments may have to consider courses of action aimed *inter alia* at reducing significantly the number of those that fall in between the cracks during such transitions: actual and potential losers are bound to become much less tolerant of their fate. This should be done *not* by attempting to slow down the pace of technical change (that would be silly and ineffectual), but on the contrary, by making sure that many more can be brought to partake in it.

### **3. From threat to promise: strategies for the AI-GPT era**

In order to meet the above mentioned challenges, governments will have to design innovative strategies in the following key areas:

1. **Education** – search for ways to provide for the changing nature of skills required for the AI era;
2. **Personal services:** these are the fastest growing occupations, but as defined at present cannot benefit from AI;
3. **Direction of technical change:** strive to human enhancing innovations, not human replacing.

#### **3.1 Education – the upcoming revolution**

As already mentioned, the expectation is that AI will become the dominant GPT of the coming era, spreading throughout the economy, and displacing in the process great many occupations. At the same time, the remaining occupations and new ones that may sprung as complementary to AI will require a new set of skills, that are not quite those currently provided by the education system, at all levels.

This is not new: the first and second Industrial revolutions in the course of the 19<sup>th</sup> century required, and were accompanied by, corresponding revolutions in education. The need to rely on a more skilled, educated workforce, as well as a more disciplined one, fed educational reforms first in Prussia (already in the late 18<sup>th</sup> century), then in the United Kingdom and in the US, that led gradually to the institutionalization of free and universal education, with highly structured, government-set curriculums.

From the late 19<sup>th</sup> century to this day, this “factory model” of education spread widely, expanding *quantitatively* in all dimensions: more hours spent at school, more subjects covered, and more years of study. Thus for example, the average years of schooling in the UK adult population was less than 1 in 1870, whereas at present it stands at over 13. Universal education now starts at age 3-4 in many countries, high school became compulsory in the second half of the 20<sup>th</sup> century, and in the past three decades some form of tertiary education has become commonplace.

It is now widely accepted that this “factory model” needs to be revised and perhaps totally revamped in view of twin pervasive developments: first, the Internet revolution, which in this context means the availability of information/knowledge on any subject, at all times and virtually at no cost; second, the rapidly changing requirements for meaningful employment.

In particular, the advent of AI as the new GPT, with its expected pervasive impact on employment, may call for a new education revolution, very much like the Industrial Revolutions of the 19<sup>th</sup> century. The key to it appears to be the shift away from imparting knowledge per se, to developing skills relevant for an AI-based economy. Likewise, such educational revolution will in all likelihood aim towards “personalized education”, departing from the quest for uniformity that has characterized education systems ever since Prussian reforms.

What are likely to be the top skills required for employment in the upcoming AI era? There is a great deal of heated discussion in this area, but some agreement is emerging around a core set of skills, such as those listed in table 1:

<b>Table 1: Skills sought for employment – from internet sites</b>		
<b>UNICEF 10 life skills</b>	<b>MyStartJob.com</b>	<b>top10onlinecolleges.org</b>
1.problem solving 2.critical thinking 3.effective communication 4.decision making 5.creative thinking 6.interpersonal relationship 7.self-awareness 8.empathy 9.coping w/stress 10.coping w/emotions	1. Communication Skills 2. Analytical & Research 3. Flexibility-Adaptability 4. Interpersonal Abilities 5. Decision making 6. Plan, Organize, Prioritize 7. Wear Multiple Hats 8. Leadership/Management 9. Attention to Detail 10. Self confidence	1.Sense Making 2.Social Intelligence 3.Novel Adaptive Thinking 4.Cross Cultural Competency 5.Computational Thinking 6.New Media Literacy 7.Transdisciplinary 8.Design Mindset 9.Manage Cognitive Load 10.Virtual Collaboration

There is a great deal of similarity between these three lists of skills, and in fact they can be classified into the following (non-exhaustive) main “types”:

- **“type I” - analytical, creative, adaptive**
  - critical & creative thinking
  - analytical & research
  - sense making
  - novel adaptive thinking
  - design mindset
- **“type II”: interpersonal, communication**
  - effective communication
  - interpersonal relationships/abilities
  - social intelligence
  - virtual collaboration
- **“type III”: emotional, self confidence**
  - self-awareness
  - empathy
  - coping w/stress
  - manage cognitive load
  - coping w/emotions

The important point to notice is that most of these skills are neither imparted in the current K-12 system, nor in academia. The whole system is still geared primarily towards the transmission of knowledge, highly structured and uniform, and not towards skills, let alone those skills. Pupils of all ages are now very aware of the fact that school-like information is available at the tip of their fingers, they are less receptive to frontal classes, their attention span is much shorter, the sort of stimuli that makes them tick is different. This is true also at the tertiary level, and in addition we are witnessing there the rise of the MOOCs (massive open online courses), and of other such online-based teaching tools.

In view of these trends, it is clear that educational strategies need to undergo equally significant changes away from the “factory system”, and the fact that the incipient GPT may render many existing occupations obsolete, provides it with renewed urgency. These are some of the issues to tackle:

- Invert the pyramid: It is now widely recognized that critical skills, hard and soft, cognitive and social, are acquired very early on. Furthermore, failure to do so at the earliest stages may be hard (even impossible) to remedy later on (see e.g. Heckman, 2014). Thus, we may have to consider investing much more in early-childhood education, from birth to age 6.
- Find ways to incorporate the development of skills (of the three types sketched above) as integral part of the teaching in every discipline and at all stages, including in academia.
- Effective educational methods are hard to come by, thus it is important to engage in bottom-up experimentation in pedagogy, school design, and social skills development, in the context of flexible, creative teaching environments.
- Reconsider the prevalent norm of uniform (typically government-mandated) curriculums and educational models, vis a vis diversity and open-innovation communities build around educational institutions.
- Foster research on the effectiveness of new educational models, their adequacy to shifting needs, the extent to which they promote equal opportunity, etc. This type of



research will be crucial given the move away from “top-down” models and the emphasis on widespread experimentation.

### 3.2 Upgrading Personal Services

A BLS study<sup>6</sup> projects that virtually all of the employment gains in the decade to 2024 will be in services, and within the service sector particularly in Health care and social assistance:

<b>Table 2</b>				
<b>US employment by major sector – millions</b>				
	<b>millions</b>			<b>Percentage growth*</b>
<b>Sector</b>	<b>2014</b>	<b>2024*</b>	<b>change*</b>	
<b>Goods producing</b>	19	19	~	~0%
<b>Services</b>	<b>121</b>	<b>130</b>	<b>+ 9.3</b>	<b>+7%</b>
<b>of which: healthcare &amp; social assistance</b>	18	<b>22</b>	<b>+ 3.8</b>	<b>+20%</b>
<b>Other</b>	10	11	+ 0.5	+1
<b>Total</b>	<b>151</b>	<b>160</b>	<b>+ 9.8</b>	<b>+6%</b>
<i>*forecast</i>				

Many of these occupations as performed today require little training and minimal educational attainment. Not surprisingly, most confer low wages, low status, and are supported by very little complementary technology. As the projections suggest, those occupations are at present not seriously threatened by AI – on the contrary, they will grow significantly. Thus, the overall prospects look rather gloomy, when not only employment is considered by also wages: major upscale occupations are projected to remain stagnant or decline, whereas low-scale occupations are expected to grow.

Is this a deterministic outcome? Not necessarily, and the case of nursing may be quite instructive. After the WWII, nursing was one of the lowest ranking occupations in the US: in 1946 the average wage of a nurse was just 1/3 that of female workers in the garment industry.<sup>7</sup> In 1964 Congress passed the *Nurse Training Act*, which essentially redefined the occupation and turned it into a profession requiring an academic degree, with an upgraded curriculum. Since then the nursing profession has risen in every dimension – salaries, status, academic requirements, range of responsibilities, etc. The nursing profession spans these days a range of specializations, whereby the upper echelon commands annual wages as high

<sup>6</sup> See: Occupational employment projections to 2024, Monthly Labor Review, U.S. Bureau of Labor Statistics, December 2015. Also in <https://www.bls.gov/opub/mlr/2015/article/occupational-employment-projections-to-2024.htm>

<sup>7</sup> In 1946, the average RN earned about one dollar an hour—or \$175 a month.

as \$100,000. Moreover, nurses use now advanced technologies, and these in turn contribute to upgrade the profession.

It could have been otherwise had it not been for the legislation of 1964, and so it is for other occupations in personal services. Thus, we need to consider proactive strategies for the *professionalization of personal services*, particularly in healthcare and education, setting standards and academic requirements.

Take for example early childhood education: in most countries there are virtually no such standards for caregivers of children aged 1-3, precisely the ages that are crucial for their development. Suppose now that they were required to have specialized academic degrees, with a curriculum that would include psychology, brain development, testing for learning disabilities, etc. Not only would the status and wages of these workers increase, but they would be much more likely to benefit from complementary advanced technologies.

The advent of AI would probably not threaten these growing occupations, and furthermore, if they were upgraded in the way just described, AI could bestow significant benefits to them as well. For that to happen smart interfaces between the practitioners of these occupations and the AI machines will have to be developed. Thus, imagine for example professional caregivers using AI to test very young children for learning disabilities, and then for treating them with specially tailored AI-based games.

To sum up: BLS projections indicate that the bulk of job creation in the decade to 2024 will be in personal services, particularly in personal care. As currently practiced, most of these occupations are at the low end of the scale, and rather impervious to technological advances. However, there are viable options to upgrade these occupations, particularly by setting academic standards and advanced curriculums. If that were to happen then the changing composition of employment (i.e. more personal care, less of many others) would not adversely affect income distribution but perhaps to the contrary; furthermore, and more importantly here, AI may play a complementary role *vis a vis* these occupations, thus raising productivity in services and triggering a virtuous cycle.

### **3.3 The Direction of Technical Change: H-enhancing or H-replacing?**

Although one of the seminal volumes in the economics of technological change is titled “The Rate and Direction of Inventive Activity”, in fact the economic discipline has traditionally dealt much more with the “rate” than with the “direction”. That may come as no surprise, since discussing “direction” requires getting into the guts of technology itself, and there is no reason to believe that economists have a comparative advantage in that regard.

Nevertheless, the extent and scope of technological advances that engulf us may require us to look more closely into the “black box”, and try to understand at the very least what types of innovations we are facing, and how they impact the economy. Furthermore, we would like to know whether there is room to affect the relative prevalence of the various types, in view of their differential economic effects.

Here is such an attempt: consider on the one hand innovations that mostly magnify, enhance and extend sensory, motoric, analytical, and other human capabilities, such as:

- **in medicine:** AI for diagnostics, e.g. for reading and interpreting x-rays, CT-scans and other imaging modalities; AI for robotic surgery (e.g. the da-Vinci robot for prostate surgery); AI-data mining of electronic medical records for follow-up evaluations of drug efficacy post-FDA approval, etc.
- **in education:** AI-based methods for “personalized teaching”; AI for online testing in MOOCs; (see also the above-mentioned applications for early childhood education), etc.

We label these “human-enhancing innovations” (“HEI”) – in medicine they don’t replace doctors but rather augment their human-bound capabilities (think of the precision and consistency of robot surgery), thus making better doctors. Similarly for teachers, eventually for judges (ruling with the aid of AI-based analysis), etc.

On the other hand, consider “human-replacing innovations” (“HRI”), that is, technical advances that *replace* human intervention, and furthermore, that often leave for humans mostly “dumb” jobs, that are not worth yet replacing given the very low wages that they command (and often are indeed difficult to replicate by machines, the proverbial one being janitors).

Some HRI’s lead to cutting-edge, virtually human-free factories (best exemplified by Telsa’s new facilities to produce batteries for its e-cars), that greatly improve productivity even if reducing employment. Consider however the polar case of Walmart, the world’s largest private employer (with over 2 million employees) – having deployed advanced technologies along its whole chain of operations, from logistics to retailing, it has turned a large proportion of its workers into “unthinking automatons”, commanding very low wages with no prospect for improvement.

These then are two types of innovations, HEI and HRI, that have very different effects on key economic and social variables. It would seem that **AI-based HEI’s** have the potential to unleash a new wave of *human creativity and productivity*, particularly in services (which to repeat are expected to be the fastest growing occupations), whereas HRI’s either decrease employment (e.g. Telsa), or create unworthy jobs.

Is it possible to design strategies to affect the direction of technical change, in the sense of stimulating HEI’s versus HRI’s? It is hard to say, but it is certainly worthwhile investigating such possibility, given the large impact that a change in direction may have on the economy. Incidentally, it would seem that in any case the traditional emphasis of economic policy on the “rate” of innovation, that is, on how much resources we devote to R&D, is misplaced – worldwide competition may be pushing us into too much investment in R&D, not too little (too many patents, too much replication, etc.). Some attention to the “direction” may bring much larger returns...

#### 4. Concluding remarks

The historical record suggests that dismal prophecies about the economic and social impact of great technological advances rarely come to pass. Thus, with AI quite poised to emerge as the new GPT, we should not necessarily envision a future whereby humans will be rendered obsolete, and mass unemployment will be the “new normal.” At the same time, as many occupations will indeed vanish, and many others will undergo significant changes, it is important to inquire into what sort of strategies may ameliorate the detrimental effects of AI, and enhance the positive ones. This is all the more important, given that in the 21<sup>st</sup> century the public at large have much less tolerance for bearing the costs of technical change, and higher expectations for sharing into its benefits here and now.

Therefore, we need to anticipate the required institutional changes, to experiment in the design of new policies, particularly in education and skills development, in the professionalization of service occupations, and in affecting the direction of technical advance. Furthermore, economists possess a vast methodological arsenal that may prove very useful for that purpose – we should not shy away from stepping into this area, since its importance for the economy cannot be overstated.

#### References

- Bresnahan, Tim and Manuel Trajtenberg, *General purpose technologies ‘Engines of growth’?* Journal of Econometrics, (65, 1), January 1995.
- Cockburn, Iain, Henderson, Rebecca and Scott Stern, *The Impact of Artificial Intelligence on Innovation*, paper presented at the AI and the Economy conference in Toronto, September 2017.
- Gordon, Robert J. 2016. *The Rise and Fall of American Growth*. Princeton: Princeton University Press.
- Heckman, James, “Fostering and Measuring Skills: Improving Cognitive and Non-Cognitive Skills to Promote Lifetime Success,” (with T. Kautz, R. Diris, B. ter Weel, and L. Borghans). Report prepared for the Organisation of Economic Co-operation and Development, Paris. (2014). <http://www.oecd.org/edu/cei/Fostering-and-Measuring-Skills-Improving-Cognitive-and-NonCognitive-Skills-to-Promote-Lifetime-Success.pdf>
- Hirschman, Albert, *Exit, Voice and Loyalty*, Harvard University Press, 1970.
- Mokyr, Joel, “The Past and the Future of Innovation: some lessons from Economic History” (2017). Mimeo, presented at this conference.
- Summers, Lawrence H. “ The Age of Secular Stagnation: What It Is and What to Do About It,” *Foreign Affairs*, Feb. 15, 2016 - <http://larrysummers.com/2016/02/17/the-age-of-secular-stagnation/>