GOSPODARKA NARODOWA

4(316)2023, 1–16
DOI: 10.33119/GN/175243
gnpje.sgh.waw.pl
SGH

The Polish Journal of Economics
Istnieje od / Published since 1931

Jakub Growiec
Department of Quantitative Economics, SGH Warsaw School of Economics, Poland

Industry 4.0? Framing the Digital Revolution and Its Long-Run Growth Consequences*

Gospodarka 4.0? Nowe ujęcie teoretyczne rewolucji cyfrowej i jej konsekwencji dla długookresowego wzrostu

Keywords:
automation, Industry 4.0, technological unemployment, Digital Revolution, long-run economic growth

JEL classification codes:
O30, O40

Article history:
submitted: 7 June, 2023
revised: 13 November, 2023
accepted: 14 November, 2023

Abstract

Are we going through a Fourth Industrial Revolution or a technological breakthrough event of an entirely different nature? In this paper, based on the hardware-software framework [Growiec, 2022; Growiec, Jabłońska, Parteka, 2023], I identify the key differences between the technologies of the Industrial Revolution (expanding our capacity to perform physical action) and the Digital Revolution (expanding our capacity to process information). I discuss the implications of these technologies for long-run economic growth, technological progress and factor demand. I find that these implications depend on the possibility of full automation of production processes, the extent of technology spillovers in R&D, and the rate of technological decay. Full automation is disruptive because it makes human labour inessential for production, potentially leading to technological unemployment as well as growth acceleration. Under positive technology spillovers in R&D, technological progress and the accumulation of R&D capital can form a dual growth engine, sustaining exponential growth even under partial automation and without population growth. As an application of the theory, I overview the effects of specific existing and hypothetical digital-era technologies, from the Jacquard loom to artificial superintelligence, for the pace of long-run growth and predicted trends in employment and factor shares.

Słowa kluczowe:
automatyzacja, gospodarka 4.0, bezrobocie technologiczne, rewolucja cyfrowa, długookresowy wzrost gospodarczy

Kody klasyfikacji JEL:
O30, O40

Historia artykułu:
nadesłany: 7 czerwca 2023 r.
poprawiony: 13 listopada 2023 r.
zaakceptowany: 14 listopada 2023 r.

Streszczenie

Czy przechodzimy właśnie przez czwartą rewolucję przemysłową czy też przełom technologiczny zupełnie innego rodzaju? W tym artykule, opierając się na modelu hardware-software [Growiec, 2022; Growiec, Jabłońska, Parteka, 2023], identyfikuję kluczowe różnice między technologiami rewolucji przemysłowej (rozszerzającymi naszą zdolność do wykonywania działań fizycznych) a technologiami rewolucji cyfrowej (rozszerzającymi naszą zdolność do przetwarzania informacji). Omawiam konsekwencje tych technologii dla długoterminowego wzrostu gospodarczego, postępu technologicznego i popytu na czynniki produkcji. Stwierdzam, że konsekwencje te zależą od możliwości pełnej automatyzacji procesów produkcyjnych, zakresu efektów ubocznych technologii w badaniach i rozwoju oraz tempa

* Financial support from the Polish National Science Center (Narodowe Centrum Nauki) under grant OPUS 19 No. 2020/37/B/HS4/01302 is gratefully acknowledged. All errors are my responsibility.

wygasania technologii. Pełna automatyzacja jest zakłócająca, ponieważ czyni pracę ludzką nieistotną dla produkcji, co może prowadzić do bezrobocia technologicznego oraz przyspieszenia wzrostu gospodarczego. Przy pozytywnych efektach ubocznych technologii w działalności badawczo-rozwojowej (B+R) postęp technologiczny i akumulacja kapitału B+R mogą stanowić podwójny motor wzrostu, utrzymujący wzrost wykładniczy nawet przy częściowej automatyzacji i bez wzrostu ludności. Jako zastosowanie wyprowadzonej teorii przedstawiam przegląd skutków rzeczywistych i hipotetycznych technologii ery cyfrowej (od krosna Jacquarda po sztuczną superinteligencję) dla tempa długoterminowego wzrostu oraz przewidywanych trendów w zatrudnieniu i udziałach czynników w tworzeniu wartości dodanej.

Introduction

Is the world currently going through a Fourth Industrial Revolution, or a technological breakthrough event of an entirely different nature? There is no doubt that growth in GDP per capita and labour productivity has massively accelerated since the beginning of the 19th century, and even more so in the 20th century (see e.g. **Galor** [**2005**], **Piketty** [**2014**], **Rosling** [**2018**]); so did technological progress. However, when it comes to identifying the key drivers of this acceleration and its technological underpinnings, the literature is divided. For some authors, such as **Gordon** [**2016**] or **Schwab** [**2016**], there is a continuity between the consecutive waves of the Industrial Revolution, see Table 1: each of them adds a new boost to the ongoing technological transformation and propels the ongoing progress. Others, however, including notably Erik Brynjolfsson and his co-authors (e.g., **Brynjolfsson, McAfee** [**2014**], **Brynjolfsson, Rock, Syverson** [**2019**]), have emphasized the qualitative difference between industrial technologies of the first two waves and digital technologies of the latter two.

Table 1. Waves of the Industrial Revolution, or the Industrial and Digital Revolution

Rev	Years	Key technologies
1.0	1770–1840	Steam engine, railroad, mechanical loom
2.0	1870–1920	Internal combustion engine, electricity, telephone, medicine, indoor plumbing
3.0	1960–2010?	Digital computer, cell phone, Internet, credit cards
4.0	2010+?	artificial intelligence, Internet of Things, cloud computing

Note: According to **Gordon** [**2016**], only waves 1–3 had a major impact on productivity; **Schwab** [**2016**] also adds wave "4.0" (and indeed largely concentrates on it). **Brynjolfsson and McAfee** [**2014**] and **Growiec** [**2022a**] emphasise the difference between waves 1–2 and 3–4 and identify the latter as a qualitatively new Digital Revolution.

Source: Author's own elaboration.

This paper subscribes to the latter view. Indeed, as I will elaborate in detail below, viewing computers, the Internet, smartphones, or artificial intelligence algorithms as achievements of the Industrial Revolution ("Industry 3.0–4.0") is oblivious to the fact that these technologies, as opposed to the technologies of "Industry 1.0–2.0", operate in the domain of information processing. Due to the non-rivalry of information [**Romer, 1990**], technologies of information processing have different properties and generate different macroeconomic implications than steam- or electricity-powered machines characteristic of Industry 1.0–2.0, expanding the humankind's capacity to perform physical action. Therefore these technologies should rather be viewed as drivers of an entirely separate technological revolution – the Digital Revolution.

It is convenient to frame this discussion using the *hardware – software framework* [**Growiec, 2022a**; **Growiec, Jabłońska, Parteka, 2023**], an approach that takes root in the basic laws of physics and, starting from there, suggests to replace the usual pair of macroeconomic factors of production – capital and labour – with two alternative aggregates: hardware ("brawn") and software ("brains"), orthogonal to the traditional distinction (Figure 1). The framework allows for reliable reduced-form modelling of production processes across both the industrial and the digital era, and offers a clear-cut distinction between *mechanization* – replacement of people with machines within hardware – and *automation* – replacement of people with machines within software.

Figure 1. Factors of production in the hardware – software framework

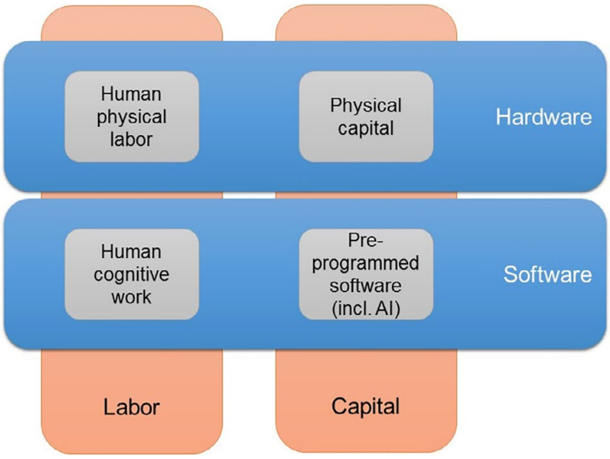

Source: Author's own elaboration.

The contribution of this paper to the literature is twofold. First, I construct a reduced form growth model with hardware and software as inputs, extending the previous findings contained in [Growiec, 2023] and allowing to identify the diverse impacts of industrial and digital technologies on long-run growth, factor shares and labour demand. Second, as an application of the theory I provide an overview of existing and hypothetical automation technologies, looking at them from three specific, theoretically motivated angles:

1. Does the technology enable full automation of complex production processes? Can production continue without any human input?
2. Does the technology contribute to positive knowledge spillovers in R&D? Does it facilitate new innovations or, to the contrary, does it make "ideas harder to find" [Bloom et al., 2020]?
3. Does the technology help reduce the rate of decay of old knowledge?

The paper links to the literature on the impacts of automation on productivity, employment, wages and factor shares [Zeira, 1998; Acemoglu, Autor, 2011; Autor, Dorn, 2013; Graetz, Michaels, 2018; Acemoglu, Restrepo, 2018a, 2018b; Andrews, Criscuolo, Gal, 2016; Arntz, Gregory, Zierahn, 2016; Frey, Osborne, 2017; Barkai, 2020; Autor et al., 2020; Jones, Kim, 2018; Hemous, Olsen, 2018]. It also addresses the debate on the macroeconomic implications of the development of AI and autonomous robots replacing humans in sophisticated, non-routine tasks requiring specialist knowledge [Hanson, 2001; Yudkowsky, 2013; Sachs, Benzell, LaGarda, 2015; Benzell et al., 2015; DeCanio, 2016; Graetz, Michaels 2018; Aghion, Jones, Jones, 2019; Berg, Buffie, Zanna, 2018; Caselli, Manning, 2019; Benzell, Brynjolfsson, 2019; Brynjolfsson, Rock, Syverson, 2019; Davidson, 2021; Trammell, Korinek, 2021; Korinek, Juelfs, 2022; Wei et al., 2022; Eloundou et al., 2023].

The remainder of the paper is structured as follows. Section 2 introduces the hardware – software framework. Section 3 derives the theoretical predictions. Section 4 discusses the implications of a range of existing and hypothetical digital-era technologies. Section 5 concludes.

Theoretical Framework

Hardware and Software as Factors of Production

The hardware – software framework [Growiec, 2022a; Growiec, Jabłońska, Parteka, 2023] presumes that all output is generated through purposefully initiated physical action. Generating output requires both some physical action involving energy – carried out by *hardware* – and some information describing the action – provided by *software*. This dichotomy mirrors the standpoint of theoretical physics, succinctly summarised by Michio Kaku: "I'm a physicist. We rank things by two parameters: energy and information."

However, both the physical action and the information processing can be carried out by either people or machines. This additional distinction underlies the map of inputs presented in Figure 1. In the hardware – software framework, physical capital and human physical labour are fundamentally substitutable inputs contributing to *hardware*; analogously, human cognitive work and pre-programmed digital software are fundamentally substitutable inputs contributing to *software*. In turn, both hardware and software are complementary and essential in the process. Furthermore, programmable hardware (computers, smartphones, robots, etc.) similarly to the human body has double duty, as a means of performing physical action and as a container for software – stored information and working algorithms.

The key building block of the hardware – software framework is the general production function featuring physical *hardware X* performing the action, and disembodied *software S* providing the information:

$$Output = \mathcal{F}\left(X,S\right). \tag{1}$$

I assume that $\mathcal{F}:\mathbb{R}_+^2 \to \mathbb{R}_+$ is increasing and concave in both factors and such that hardware X and software S are essential (i.e., $\mathcal{F}(0,S)=\mathcal{F}(X,0)=0$) and mutually complementary in production (the elasticity of substitution between X and S is below unity).

Hardware X ("brawn") encompasses both the services of physical capital K and human physical labour L, where the latter variable excludes any know-how or skill of the worker. In turn, software S ("brains") includes the skills and technological knowledge employed in human cognitive work, H, as well as pre-programmed software Ψ – instructions to be performed by programmable hardware. Pre-programmed software Ψ includes AI algorithms that learn and improve their performance with new data.

Within hardware X, capital and labour are inessential and substitutable as agents of physical action (elasticity of substitution above unity). This reflects the idea that the same physical actions should lead to the same outcomes, regardless of whether they are performed by a human or machine. Analogously with software S: if instructions are the same, then the outcomes should be the same too, regardless of whether they have been generated by a human brain or a digital information processing unit.

However, we do not know yet whether all essential tasks will eventually be automatable in the future. This is important because if certain essential cognitive tasks are never automated, then the reduced-form production function should rather assume gross complementarity (elasticity of substitution below unity) between human cognitive work and pre-programmed software within software S. For a comprehensive discussion of the mapping of partial vs. full task automatability to the elasticity of substitution, please refer to [**Growiec, 2022b**].

These assumptions lead to the following specification of hardware and software:

$$X = G_1(L,K), \;\; S = G_2(H,\Psi), \tag{2}$$

where the elasticity of substitution in G_1 is above one, and in G_2 – above one in the full automation scenario, and below one in the partial automation scenario. The functions $G_1:\mathbb{R}_+^2 \to \mathbb{R}_+$ and $G_2:\mathbb{R}_+^2 \to \mathbb{R}_+$ are assumed increasing and concave in both inputs. The replacement of L with K within hardware X will be referred to as *mechanisation*, whereas the replacement of H with Ψ within software S will be called *automation*.

In line with the economic interpretation, human physical labour $L=\zeta N$ is assumed to be rivalrous and proportional to the number of workers, $L=\zeta N$. Physical capital K is also rivalrous but can be unboundedly accumulated in per capita terms. In turn, human cognitive work $H = AhN$ is proportional to technological knowledge A, skill level h, and the number of workers N. Technological knowledge A, also interpreted as the size of the repository of codes, is non-rivalrous [**Romer, 1986**; **1990**] and accumulable. Skill-adjusted cognitive labour is rivalrous though. Finally, pre-programmed software $\Psi = A\psi\chi K$ is assumed to be proportional to technological knowledge A (as above), algorithmic skill level ψ and the stock of programmable hardware χK on which the software is run. It is assumed that pre-programmed software scales with programmable hardware χK because of the near zero cost of copying digital software.

Finally, the hardware – software framework views technological progress (growth in A) as expansion of the repository of codes. It is the *informational* character of technological knowledge A which makes it non-rivalrous and, therefore, a source of increasing returns to scale [Romer, 1986; 1990]. Unlike in Romer's original contributions, however, in the hardware – software framework technological knowledge can be put to productive use by both humans and machines. As a result, technological progress is modeled as *software-augmenting*.

With these assumptions in hand, the general production function takes the form:

$$Output = \mathcal{F}\left(G_1\left(\zeta N, K\right), G_2\left(AhN, A\psi\chi K\right)\right). \tag{3}$$

Reduced-Form Economic Growth Model

Hardware and software are inputs of both *production* and *R&D*. The output of production is the final good Y, serving both consumption and investment purposes, whereas R&D produces new technological ideas $\dot{A}$. In what follows, I modify the framework considered in Growiec [2023], Growiec, Jabłońska, and Parteka [2023] to discuss a range of important new scenarios that have not been yet analysed, but which may be critical for the assessment of the characteristics and potential of certain existing or hypothetical digital technologies.

I consider the following reduced-form two-sector economic growth model with a production and R&D sector:

$$\dot{A} = \Phi\left(G_1\left(\zeta N, K\right), G_2\left(A^{\phi} hN, A^{\phi}\psi\chi K\right)\right) - \omega A, \tag{4}$$

$$Y = F\left(G_1\left(\zeta N, K\right), G_2\left(AhN, A\psi\chi K\right)\right), \tag{5}$$

$$\dot{K} = sY - \delta K, \tag{6}$$

where the term A^{ϕ} (with $\phi \in \mathbb{R}$) captures *software-augmenting knowledge spillovers* in R&D. These knowledge spillovers represent either "standing on shoulders" effects if $\phi > 0$ [Jones, 1995] or "fishing out ideas" effects if $\phi < 0$ [Bloom et al., 2020]. In turn, $\omega \geq 0$ captures the rate of knowledge decay, interpreted also as knowledge depreciation or obsolescence.

The aggregate production function $F: \mathbb{R}^2_+ \to \mathbb{R}_+$ and the idea production function $\Phi: \mathbb{R}^2_+ \to \mathbb{R}_+$ are assumed to have the properties of the general production function $\mathcal{F}$ discussed above.[1] I additionally assume that F, G_1, G_2 and Φ are characterised by constant returns to scale. Time $t \geq 0$ is continuous.

Finally, I posit that bounded variables (s, h, ψ, χ) will eventually stabilise, and so will global human population N. Therefore in my derivations I concentrate solely on the dynamics of two state variables of the model, K and A, in the long-run limit, treating s, h, ψ, χ and N as given constants – just like in the canonical Solow model (cf. Solow [1956], Jones [2005]). I also abstract from the optimal allocation of factors and instead assume that fractions of hardware and software used in production vs. R&D are always constant. They are included in functions F, Φ for compact notation. Analogously, fractions of capital and labour used in hardware vs. software are also assumed constant and are included in functions G_1, G_2.

The two key differences between the above model and that of Growiec [2023] are that (i) knowledge spillovers in R&D are now assumed to be software-augmenting, and (ii) knowledge is subject to gradual decay over time.

As regards knowledge spillovers in R&D, the assumed specification (5) should be contrasted against two literature benchmarks. First, if $\phi = 1$ in (5), then the stock of technological knowledge would be included in the software term symmetrically in production and R&D [Romer, 1990]. In comparison to this benchmark, $\phi > 1$ implies the existence of positive knowledge spillovers, contributing positively to R&D output beyond the contribution of software S itself. Then R&D output grows more than proportionally with A. In turn, $\phi < 1$ implies

1 Using a CES ideas production function specification, [Growiec, McAdam, Mućk, 2023] demonstrate empirically for the United States that capital and labour are gross complements also in R&D ($\sigma_{R\&D} < 1$).

the existence of negative knowledge spillovers, diminishing the contribution of software S to R&D output. Then R&D output grows less than proportionally with A. Finally, if $\phi < 0$, then we arrive at a scenario where the negative knowledge spillovers are so strong that they imply "fishing out effects": ideas are getting harder to find [**Bloom et al., 2020**], and R&D output *declines* with A.

The second literature benchmark is due to **Jones** [**1995**]. In this seminal paper, as well as in all R&D-based growth literature it spurred, knowledge spillovers in R&D were viewed as "Hicks-neutral," i.e. affecting R&D output multiplicatively. This exact approach has also been taken in the related paper by **Growiec** [**2023**]. However, the theoretical base of the hardware – software framework suggests that knowledge spillovers in R&D should be modelled as software-augmenting rather than Hicks-neutral. The reason is that Hicks neutrality implies that technical change is also to some degree hardware-augmenting. However, hardware-augmenting technical change, interpreted as improvements in the energy efficiency of physical actions, cannot be *unbounded* – as required in my analysis based on (5) – because that would contradict the laws of thermodynamics [**Beaudreau, Lightfoot, 2015**].

As regards the gradual decay of knowledge, the term $-\omega A$ is usually disregarded in the growth literature; technology improvements are assumed eternal. However, while the assumption $\omega = 0$ is a useful approximation facilitating analytical tractability, it is not necessarily guaranteed to hold in reality. First, there are known historical episodes when scientific knowledge was lost, such as following the fall of the ancient Roman empire. Second, technological knowledge may be subject to decay as science progresses, mirroring the obsolescence of old knowledge in light of new findings. Third, mounting technological knowledge may gradually become harder to access given the cognitive limitations of the human brain [**Jones, 2009**; **Hidalgo, 2015**; **Growiec, 2022a**].

All these caveats notwithstanding, the rate of knowledge decay is probably very small these days, after it has fallen over centuries thanks to discoveries which dramatically increased the durability of knowledge, such as the alphabet, writing, the printing press (before the Industrial Revolution), telecommunications and audiovisual storage of data (in the industrial era), and ultimately digital memory and the Internet with efficient search engines (in the digital era).

Basic Asymptotic Results

I obtain the following approximations as $K/N \to \infty$ (cf. **Growiec** [**2023**]):

$$X \approx \alpha K \text{ where } \alpha = G_1(0,1) = \lim_{x\to 0} G_1(x,1), \tag{7}$$

$$S \approx \beta A\psi\chi K \text{ where } \beta = G_2(0,1) = \lim_{x\to 0} G_2(x,1), \text{ full automation,} \tag{8}$$

$$S \approx \gamma AhN \text{ where } \gamma = G_2(1,\infty) = \lim_{y\to \infty} G_2(1,y), \text{ partial automation.} \tag{9}$$

These approximations are valid thanks to the assumption of constant returns to scale as well as above unitary elasticity of substitution in G_1 (equation 7) and in G_2 (equation 8), and alternatively, below unitary elasticity of substitution in G_2 (equation 9).

I also use the following asymptotic notation:

$$a_K = F(1,\infty) = \lim_{y\to\infty} F(1,y), \; b_K = \Phi(1,\infty) = \lim_{y\to\infty} \Phi(1,y), \tag{10}$$

$$a_N = F(\infty,1) = \lim_{x\to\infty} F(x,1), \; b_N = \Phi(\infty,1) = \lim_{x\to\infty} \Phi(x,1). \tag{11}$$

By the assumptions of constant returns to scale and less than unitary elasticity of substitution in F and Φ, limits in (10)–(11) exist and are finite.

Predictions for the Digital Era

The reduced-form growth model (4)–(6) produces the following long-run predictions.

Scenario #1: Full Automation of Production and R&D, $\phi > 0$

In this scenario, as $K / N \to \infty$, for computing the long-run dynamics we may approximate $X = G_1(\zeta N, K) \approx \approx \alpha K$ and $S = G_2(AhN, A\psi\chi K) \approx \beta A\psi\chi K$.

Assuming $\phi > 0$, I obtain the following asymptotic results as $K \to \infty$ and $A \to \infty$:

$$Y = F(\alpha K, \beta A\psi\chi K) = \alpha K F(1, A\psi\chi \cdot \beta / \alpha) \to \alpha a_K K, \tag{12}$$

$$\dot{A} = \alpha K \Phi(1, A^{\phi}\psi\chi \cdot \beta / \alpha) - \omega A \to \alpha b_K K - \omega A, \tag{13}$$

$$\dot{K} \approx (s\alpha a_K - \delta) K. \tag{14}$$

Hence, when all essential tasks are automatable, the output growth rate will converge to $g = g_K = g_A = = s\alpha a_K - \delta$, and the long-run growth engine will be the accumulation of programmable hardware [**Jones, Manuelli, 1990**].

The full automation scenario presumes that all essential tasks, both physical and cognitive, will be eventually performed by machines. This follows from the gross substitutability of people and machines within both hardware and software, coupled with the unbounded accumulation of machines per worker (K / N). Furthermore, with "standing on shoulders" effects in R&D ($\phi > 0$, covering both the case of positive knowledge spillovers if $\phi > 1$ and weak negative knowledge spillovers if $\phi \in (0, 1)$), software will be growing systematically faster than hardware, so that asymptotically, given the complementarity between hardware and software, the role of software-augmenting R&D as a source of growth will fall to zero. Accordingly, the software share of output and the aggregate labour share will gradually decline toward zero as well. In this scenario, demand for labour declines, indicating a risk of technological unemployment unless people would be willing to work for any wage [**Growiec, 2022b**; **Korinek, Juelfs, 2022**]. At the same time, as full automation makes it possible to get rid of the development bottleneck due to the limited supply of human cognitive work, this scenario offers the prospect of a substantial growth acceleration.

Scenario #2: Full Automation of Production and R&D, $\phi \leq 0$

The only difference between the current and the previous scenario is that knowledge spillovers in R&D are now strongly negative (there are "fishing out" effects, $\phi \leq 0$). Under full automation, this changes the implications only slightly.

Assuming $\phi \leq 0$, I find that asymptotically as $K \to \infty$ and $A \to \infty$:

$$Y = F(\alpha K, \beta A\psi\chi K) = \alpha K F(1, A\psi\chi \cdot \beta / \alpha) \to \alpha a_K K, \tag{15}$$

$$\dot{A} = \beta\psi\chi A^{\phi} K \Phi\left(\frac{\alpha}{\beta\psi\chi A^{\phi}}, 1\right) - \omega A \to \beta b_N \psi\chi A^{\phi} K - \omega A, \tag{16}$$

$$\dot{K} \approx (s\alpha a_K - \delta) K. \tag{17}$$

Hence, the output growth rate will again converge to $g = g_K = s\alpha a_K - \delta$, and the long-run growth engine will be again the accumulation of programmable hardware. Ideas A will follow suit, but due to "fishing out" effects in R&D, their growth will be slower in this scenario: $g_A = \frac{g}{1-\phi} \leq g$.[2]

2 Proof by solving a Bernoulli differential equation.

With strongly negative knowledge spillovers ("fishing out" effects), the software input to R&D will be growing systematically slower than hardware. Asymptotically, all growth will be generated by the accumulation of machines, owing to the asymptotically linear production function. Software-augmenting R&D will provide a negative contribution to growth, albeit the drag from this factor will be bounded and will eventually fall to zero. Accordingly, the software share of output and the labour share will gradually decline toward zero, and so will the demand for labour, indicating a threat of technological unemployment. As full automation makes it possible to get rid of the development bottleneck due to the limited supply of human cognitive work, this scenario offers the prospect of a substantial growth acceleration as well.

Scenario #3: Partial or No Automation, $\phi > 1$

As $K/N \to \infty$, for computing the long-run dynamics we may approximate $X = G_1(\zeta N, K) \approx \alpha K$ and $S = G_2(AhN, A\psi\chi K) \approx \gamma AhN$.

The "partial or no automation" scenario presumes that cognitive tasks cannot be fully automated, and therefore people and machines will always remain gross complements within software. This renders human cognitive work an essential and asymptotically limiting factor of production.

If $\phi > 1$ then an asymptotic balanced growth path is attained as $K \to \infty$ and $A \to \infty$, with

$$g = g_A = \alpha b_N \frac{K}{A} - \omega = sF\left(\alpha, \gamma \frac{A}{K} hN\right) - \delta. \tag{18}$$

With positive knowledge spillovers ($\phi > 1$) the effective supply of software in the R&D sector will be growing faster than the supply of hardware (R&D capital). Therefore R&D capital will eventually become the bottleneck of long-run growth. In the long run, the output growth rate is determined by the pace of R&D, which is in turn sustained by the accumulation of R&D capital. This is a *dual growth engine*, and both hardware and software grow asymptotically at the same rate g [**Growiec, 2022c**], which depends on the parametrisation of the aggregate production function F. Accordingly, the software share of output and the labour share are expected to stabilise at some interior level between zero and one. The demand for labour will be sustained, fending off the risk of technological unemployment and guaranteeing that wages will grow at the same rate as output. In this scenario, there is no clear indication of growth acceleration or deceleration compared to the industrial era.

Scenario #4: Partial or No Automation, $\phi = 1$

If there are no knowledge spillovers in R&D ($\phi = 1$), then an asymptotic balanced growth path is attained as $K \to \infty$ and $A \to \infty$, with

$$g = g_A = \Phi\left(\alpha \frac{K}{A}, \gamma hN\right) - \omega = sF\left(\alpha, \gamma \frac{A}{K} hN\right) - \delta. \tag{19}$$

In this case, the long-run output growth rate is determined by the pace of R&D, which is in turn sustained by the accumulation of R&D capital. This is again a dual growth engine, and both hardware and software grow asymptotically at the same rate g [**Growiec, 2022c**]. The only difference with respect to the previous scenario is that without knowledge spillovers, both R&D capital K and software S (proportional to AN) are now growth bottlenecks, and accordingly the long-run growth rate depends on both the parametrisation of the aggregate production function F and the ideas production function Φ. The conclusions for factor shares and labour demand remain intact. Again, there is no clear indication of growth acceleration or deceleration compared to the industrial era.

Scenario #5: Partial or No Automation, $\phi < 1, \omega = 0$

At this point, it is helpful to note that none of the above discussions mentioned the rate of knowledge decay ω. This is because in all earlier scenarios, growth engines related to the accumulation of programmable hardware and, in some cases, also software-augmenting technical change, were sufficiently strong to outrun knowledge decay. However, this is no longer guaranteed in the case of partial or no automation and negative knowledge spillovers ($\phi < 1$). I will first consider the case without knowledge decay ($\omega = 0$), and then proceed to the case with knowledge decay ($\omega > 1$).

With $K / N \to \infty$, for computing the long-run dynamics we may again approximate $X = G_1(\zeta N, K) \approx \alpha K$ and $S = G_2(AhN, A\psi\chi K) \approx \gamma AhN$. If $\phi < 1$ and $\omega = 0$ then asymptotically there is unbounded but less than exponential growth as $K \to \infty$ and $A \to \infty$ [**Groth, Koch, Steger, 2010**; **Philippon, 2022**], with

$$Y = F(\alpha K, \gamma AhN) \to \gamma a_N AhN, \tag{20}$$

$$\dot{A} = \Phi(\alpha K, \gamma A^{\phi} hN) \to \gamma b_N A^{\phi} hN. \tag{21}$$

In this case, the long-run output growth rate is determined by the pace of software-augmenting R&D. Solving (21), I find that A grows as a power function of time, $A(t) \propto t^{\frac{1}{1-\phi}}$. Capital also grows unboundedly, and also sub-exponentially.

In this scenario, only the effective supply of software in the R&D sector is a growth bottleneck. The software share of output and the aggregate labour share will grow towards unity. The demand for labour, specifically cognitive work in the R&D sector, will be systematically increasing. There will be no technological unemployment, but there will be a secular growth slowdown.

Scenario #6: Partial or No Automation, $\phi < 1$, $\omega > 0$

In this scenario, there is least growth potential as neither full automation nor technological progress can sustain growth over the long run. Specifically the latter growth engine, with negative knowledge spillovers, finds it impossible to outrun knowledge decay. I find that if $\phi < 1$ and $\omega > 0$ then the economy converges to a steady state, satisfying

$$0 = \dot{K} = sF(\alpha K, \gamma AhN) - \delta K, \tag{22}$$

$$0 = \dot{A} = \Phi(\alpha K, \gamma A^{\phi} hN) - \omega A. \tag{23}$$

Hence, growth eventually comes to a stop. The crucial growth bottleneck is located in the R&D equation, reflecting the scarcity of software and specifically ideas for producing new ideas. All factor shares stabilise at some interior values.

Taking Stock

We are now in a position to take stock of the results obtained so far. As summarised in Table 2, the future of economic growth in the ensuing digital era depends on (i) the possibility of full automation, (ii) whether knowledge spillovers in R&D are positive or negative, and (iii) whether knowledge is subject to decay. Premises (i)–(iii) have been presented in decreasing order of importance: specifically if (i) full automation is possible, then irrespective of (ii)–(iii) one should expect an acceleration of growth, driven by the accumulation of programmable hardware, a decline of the labour share towards zero, and potential for technological unemployment (unless people would work for any wage). Otherwise, if (i) full automation is not possible, then there is no growth acceleration and (ii) the magnitude of knowledge spillovers in R&D becomes key. Under positive or no spillovers, exponential growth is maintained thanks to a dual growth engine: R&D plus the accumulation of R&D capital. The labour share then stabilises at an interior value. Otherwise, if (ii) knowledge spillovers in R&D are negative, then exponential growth cannot be sustained, and then it matters whether (iii) there is knowledge decay. Without knowledge decay, there is less-than-exponential but unbounded growth, and the key growth engine is R&D, which produces cognitive labour-augmenting technological progress. The labour share goes up towards unity. In turn, with knowledge decay, growth peters out and the economy converges to a steady state.

Table 2. Summary of theoretical results

Scenario	Growth engine	Growth rate
Full Automation, $\phi>0$	K acc	$g=s\alpha a_K-\delta, g_A=g$
Full Automation, $\phi\leq 0$	K acc	$g=s\alpha a_K-\delta, g_A<g$
Partial/No Automation, $\phi>1$	K acc+LATC	$g_A=g$, depends on F (18)
Partial/No Automation, $\phi=1$	K acc+LATC	$g_A=g$, depends on F,Φ (19)
Partial/No Automation, $\phi<1, \omega=0$	LATC	less than exponential
Partial/No Automation, $\phi<1$, $\omega>0$	LATC	steady state

Notes: K acc – accumulation of programmable hardware; LATC – cognitive labour-augmenting technical change.
Source: Author's own elaboration.

Overview of Technologies

Let me now present the potential of a range of existing and hypothetical technologies against the theoretical backdrop of the hardware – software framework.

Under this framing, technological progress can be divided into progress in humankind's capacity to (a) perform physical action and (b) process information. Therefore the foremost distinction to observe is between *mechanisation* and *automation* technologies, replacing people with machines within *hardware* and *software* respectively. Of course, no innovation implies a direct one-to-one substitution. Instead, thanks to the new machines, the overall productive capacity goes up, and accordingly the productivity of complementary inputs goes up. For example, the excavator is a mechanisation technology that replaces the muscle power of a man with a shovel with the power of the excavator's engine, while simultaneously increasing the overall digging capacity per unit of time and augmenting the cognitive work of the person operating the excavator, who is now able to work much more productively. By contrast, Microsoft Excel is an automation technology that replaces the cognitive work of an office clerk or a company manager. In its essence, it replaces mathematical operations carried out in the human brain with digital operations carried out in the computer, while simultaneously increasing the overall computational capacity per unit of time and augmenting the unautomated higher-order cognitive tasks, such as managing the finances of a firm.

Mechanisation Technologies

Mechanisation technologies were the cornerstone of the first two waves of the Industrial Revolution – Industry 1.0–2.0. While there were some accompanying developments during the 19th century, which also greatly helped humankind, such as progress in medicine and hygiene, indoor plumbing or the telephone, the key drivers of progress at that time were technologies that harnessed the energy encapsulated in fossil fuels and put it into productive use.

The key invention of Industry 1.0 was the steam engine, which harnessed the power of hot steam generated by burning coal. The steam engine was subsequently applied in pumps, steel mills, rail transport, steamboats, and further in a variety of industrial applications, such as in the cotton industry in 19th-century England. Subsequently, Industry 2.0 added the internal combustion engine and the electric engine. Electricity proved to be a particularly potent vehicle of technological progress, with a plethora of applications that became possible thanks to the separation of electricity generation (in power plants) from end use. Electrification provided access to power across all businesses as well as households.

Mechanisation technologies provided a massive contribution to economic growth in the 19th and 20th centuries by increasing the overall productive capacity, specifically by increasing the available physical power and the speed of physical actions. Their potential is best realised when the machines performing the requisite physical actions have instant access to an energy source and physical action is generated without the input of human physical work. Although new mechanisation technologies are developed to this day, the *economies of speed* have already been largely exhausted [**Beaudreau, 2020**]. Moreover, the available power is constrained by the mounting environmental toll from burning fossil fuels – at least until we manage to switch to clean renewable energy.

Mechanisation technologies operate in the domain of hardware, and it seems uncontroversial to assume that all production processes can be fully mechanised, making human physical work inessential. There is also no upper bound in sight on the accumulation of physical capital per worker. What is key, though, is whether this growth potential can be effectively tapped by the complementary software. It will be the automation technologies, not mechanisation technologies, that can reduce this scarcity and thus play the key role in determining our future in the digital era.

Automation Technologies

Automation technologies used for information processing are key to what is sometimes called "Industry 3.0–4.0". They had some earlier predecessors, though. Certain mechanical devices, such as the ancient Antikythera mechanism or medieval tower clocks, processed information in an autonomous way for centuries. Next, during the first wave of the Industrial Revolution, a notable breakthrough was achieved with the Jacquard loom, the first *programmable* mechanical device which used punch-card coding. A few decades later, another set of 19^{th}-century inventions facilitated the transmission of information: the telegraph and telephone.

But it was not until the digital computer, developed immediately after World War II, that automation entered the economy and society on a larger scale. Following the arrival of the computer, the next key developments of the Digital Revolution were the Internet, cellular telephony, broadband, and smartphone. The celebrated recent achievements of "Industry 4.0", or the second wave of the Digital Revolution, include the Internet of Things, cloud computing, and artificial intelligence (AI) algorithms [**Schwab, 2016**].

Automation technologies contribute to economic growth by accelerating the flow of information, improving the market allocation, and chiefly by increasing the cumulated computing power of humankind. Their potential is best realised when the machines performing the requisite computations have instant access to information sources (code and data), and information is processed without the input of human cognitive work. Furthermore, as information is non-rivalrous [**Romer, 1990**], automation technologies contribute to economic growth also by exploiting increasing returns to scale.

Some of the growth effects of automation technologies are achieved from merging hardware and software. Examples of such technologies are *autonomous machines* such as industrial robots, household robots, self-driving vehicles, and smart appliances, as well as *computer – human interfaces*, from old-fashioned keyboards and monitors to futuristic augmented/virtual reality devices.

Characteristics of Automation Technologies

Let me now characterise automation technologies, focusing on the three key questions emphasised by the theory.

1. Does the technology enable full automation of production processes? Can production continue without any human input?

Thus far we can only partially automate production processes, and our automation technologies are applied primarily to routine, repetitive and easily codifiable tasks [**Acemoglu, Autor, 2011**; **Autor, Dorn, 2013**; **Frey, Osborne, 2017**]. While the set of automatable tasks is gradually expanding, the key advantage of the human brain compared to digital-era machines lies firmly with our broad and adaptive intelligence [**Bostrom, 2014**; **Brynjolfsson, McAfee, 2014**]. Automation will always be only partial as long as machines do not improve upon the breadth of their applicability (*the same* software algorithm must be able to perform a wide range of tasks) and adaptation capability (the algorithm must also be able to *learn* to perform a variety of new tasks).

It appears that there is just one hypothetical technology that could be a true game changer: development of adaptive, general-purpose AI, potentially culminating in superhuman artificial general intelligence (AGI) with the capability to self-improve and self-replicate. Setting aside the multitude of threats such hypothetical technology would generate (cf. **Yudkowsky** [**2013**], **Bostrom** [**2014**], **Ord** [**2020**], **Aschenbrenner** [**2020**], **Trammell** [**2021**], **Growiec** [**2022a**]), it seems to be the only possible way towards full automation. Otherwise, complex production processes would always be only partly automated, and the unautomated essential tasks performed by people would constitute a growth bottleneck, making growth acceleration impossible while keeping firm demand for human cognitive work.[3]

However, while the future of the digital era will be crucially determined by the progress in the breadth and adaptivity of AI, there is also a range of auxiliary technologies needed for the AI algorithms to effectively operate: general-purpose computing power, digital memory, the Internet, bandwidth, provision of data (using e.g. sensing technologies), and reliable access to energy.

2. Does the technology contribute to positive knowledge spillovers in R&D? Does it facilitate new innovations or, to the contrary, does it make "ideas harder to find"?

   Many digital technologies are contributing to increase the extent of positive knowledge spillovers in R&D. Specifically, digitisation of the knowledge base (datasets, books, journals, patents, blueprints, etc.), coupled with fast data transmission over the Internet and the development of efficient web search engines, helps organise the knowledge and efficiently access the existing information. Digital technologies also facilitate cooperation in geographically dispersed research teams. At the same time, partial automation of R&D tasks, thanks to the use of computers with research software, including, to an increasing extent, AI algorithms, helps researchers exploit some long-distance links in the knowledge space. These links were previously unexplored because of the cognitive limitations of the unaugmented human brain, resulting in specialisation in increasingly narrow fields of scientific inquiry.

   Meanwhile, physical limits to progress may exist in specialised lines of research, as seen in economies of speed or digital computing capacity measured by the number of transistors in an integrated circuit. The existence of such limits implies that, barring a new breakthrough, new ideas may get harder to find over time [**Bloom et al., 2020**]. However, this regularity is general and not affected by the Digital Revolution. Thus, on balance, one should expect that knowledge spillovers have improved during the Digital Revolution.

3. Does the technology reduce the rate of decay of old knowledge?

   In principle, the digitisation of the knowledge base and its storage on a decentralised Internet should work to reduce the rate of decay of old knowledge. It does incur the risk of the knowledge being destroyed or maliciously manipulated, but this risk can be largely mitigated by responsible backup policies. However, one could argue that the rate of knowledge decay is very low, perhaps close to zero, not just in the digital era, but since the popularisation of printed books during the Renaissance period.

[3] Google Chief Economist Hal Varian advises: "Seek to be a scarce complement to increasingly abundant inputs."

Overview of Hypothetical Automation Technologies

Now I would like to present an overview of hypothetical automation technologies which can be potentially developed in the future.

The first group of possible developments includes autonomous machines specialised in certain pre-defined tasks, such as self-driving vehicles, care robots, cooking robots, and smart refrigerators renewing their inventory. None of these should not have an impact on the three key questions discussed above.

The second group includes new programmable hardware, allowing significant increases in computation capacity, such as quantum computing or biological computing. Humankind may one day develop practically useful computers that are no longer based on semi-conductors and integrated circuits, possibly increasing the available cumulative computing power by orders of magnitude. This would massively boost further progress in digital technologies, but on its own – without a sufficient improvement in the quality of code – it would act just like the development of new sources of energy, and thus should not have any impact on the three key questions discussed above.

The third group includes technologies enhancing the interconnectedness between people and machines, such as virtual/augmented reality devices, implantable brain – machine interfaces, healing nanorobots, etc. These technologies should have the potential to improve knowledge spillovers in R&D by improving the cognitive capacities of researchers. However, enhancing the capabilities of the human mind in our "race against the machine," if anything, may slow down the potential transition to full automation.

The final group includes AI algorithms. These are, for example, generative AI algorithms producing text, visual arts and music, e.g., large language models such as GPT-4 [**OpenAI, 2023**; **Korinek, 2023**]; web search engines, such as Google; and a variety of other algorithms producing predictions based on big datasets, such as AlphaZero and AlphaFold. As argued above, AI algorithms have the potential of driving a disruptive shift from partial to full automation. They are also conducive to the creation of positive knowledge spillovers in R&D, both via search engines and helping guide human intuition [**Davies et al., 2021**]. Moreover, AI algorithms can have the ability to learn, self-improve (rewrite their own code) and self-replicate. They may potentially also tap into the computational capacity of quantum or biological computers. For all these reasons, this last group of developments appears most important to watch in the future.

In fact, superhuman AGI may arrive quite soon. We are currently observing a race between OpenAI (GPT), Google (Gemini, LaMDA, PaLM), Anthropic (Claude) and Meta (LLaMA) towards ever more capable and generally applicable AI. The AI capabilities frontier is advancing at breakneck speed. Each week new developments are announced, increasing both the optimisation power of state-of-the-art AI and the breadth of its application. The timelines to superhuman AGI, while subject to massive uncertainty, shortened significantly after the rollout of GPT-4, with the current (November 2023) median prediction of AI experts at metaculus.com being as early as 2032. OpenAI themselves provide an even sharper timeline, suggesting that AGI will be created most likely before 2030: "While superintelligence seems far off now, we believe it could arrive this decade."

Concluding Remarks

In this paper, I have constructed a reduced form growth model with hardware and software as inputs, making it possible to identify the diverse impacts of industrial and digital technologies on long-run growth, factor shares and labour demand. As an application of the theory, I have also provided an overview of existing and hypothetical automation technologies, looking at them from three specific, theoretically motivated angles:

1. Does the technology enable full automation of complex production processes? Can production continue without any human input?
2. Does the technology contribute to positive knowledge spillovers in R&D? Does it facilitate new innovations or, to the contrary, does it make "ideas harder to find"?
3. Does the technology help reduce the rate of decay of old knowledge?

The theoretical analysis suggests that the first question is key as knowledge spillovers in R&D and the rate of knowledge decay become essentially irrelevant under full automation. Accordingly, in my overview of automation technologies, I have emphasized that full automation can be achieved in the future only through the development of increasingly powerful, adaptive, general-purpose AI.

When presented against the backdrop of the long history of human civilisation [**Growiec, 2022a**], the Industrial and Digital Revolutions constitute the two most recent swings of the great pendulum of economic development (Figure 2), following the earlier Cognitive Revolution (ca. 70,000 years ago – a revolution in learning), Agricultural Revolution (ca. 10,000 years ago – in access to energy), and the Scientific Revolution (around 1550 CE – in learning). This perspective further underscores that the Industrial Revolution, which dramatically improved humankind's access to energy, should be considered separately from the Digital Revolution, which dramatically improved learning by decoupling humankind's overall information processing power from the cognitive capabilities of the human brain. By contrast, terms such as "Industry 3.0–4.0," which ignore this distinction, are misleading.

Figure 2. The great pendulum of economic development

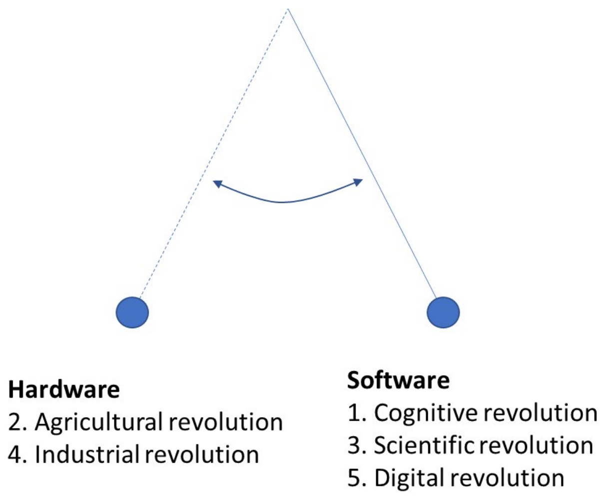

Source: [**Growiec, 2022a**]; hardware revolutions are breakthroughs in humankind's access to energy; software revolutions are breakthroughs in humankind's ability to process information and learn.

The great pendulum of economic development swings back and forth between hardware and software, at each point in time highlighting the factor which at a specific time undergoes relatively faster expansion. Its driving force is the human drive to alleviate the scarcity of resources [**Growiec, 2022a**]. Each subsequent technological revolution is achieved when a mounting scarcity of a factor, whether hardware or software, is resolved thanks to a major technological breakthrough. Specifically, in the industrial era, thanks to the rapid accumulation of physical capital in the form of machines using energy from fossil fuels, humankind managed to overcome the energy scarcity, and eventually information has become relatively scarce, prompting an education boom, creation of more inclusive and democratic institutions, and finally the Digital Revolution. Currently, in the ongoing digital era, we are witnessing a massive buildup of digital data processing power, data volume and bandwidth [**Hilbert, López, 2011**], which helps overcome the information bottleneck. If this progress continues unabated, particularly in the field of AI, energy may eventually become relatively scarce again, prompting the need for the next technological revolution in hardware.

Needless to say, this perspective comes with massive caveats. First and foremost, the prospects of developing ever more powerful, adaptive, general-purpose AI raises the question of human control. Autonomous machines, once installed, must be granted discretion in their decision making; over time, the potential impact of these decisions will only grow. The question then is, how much autonomy are people ready to give up, sacrificing their control locally with the hope of achieving accelerated growth and greater autonomy in the bigger picture. Secondly, at the end of the line, the hypothetical superintelligent artificial general intelligence (AGI) may contain a maximisation routine that could become a runaway process, escaping any human control and posing an existential threat to humanity [**Bostrom, 2014**; **Ord, 2020**].

References

Acemoglu D., Autor D. [2011], Skills, Tasks and Technologies: Implications for Employment and Earnings, in: Ashenfelter O., Card D. (eds.), *Handbook of Labor Economics*: 1043–1171, Elsevier.

Acemoglu D., Restrepo P. [2018a], The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment, *American Economic Review*, 108: 1488–1542.

Acemoglu D., Restrepo P. [2018b], Modeling Automation, *AEA Papers and Proceedings*, 108: 48–53.

Aghion P., Jones B.F., Jones C.I. [2019], Artificial Intelligence and Economic Growth, in: Agrawal A., Gans J., Goldfarb A. (eds.), *The Economics of Artificial Intelligence: An Agenda:* 237–282, University of Chicago Press.

Andrews D., Criscuolo Ch., Gal P.N. [2016], *The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective*, Working Party No. 1 on Macroeconomic and Structural Policy Analysis, OECD.

Arntz M., Gregory T., Zierahn U. [2016], *The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis*, OECD Social, Employment and Migration Working Paper No. 189, OECD Publishing, Paris.

Aschenbrenner L. [2020], *Existential Risk and Growth*, Global Priorities Institute WP 6–2020, University of Oxford.

Autor D.H., Dorn D. [2013], The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market, *American Economic Review*, 103(5): 1553–1597.

Autor D., Dorn D., Katz L.F., Patterson Ch., Van Reenen J. [2020], The Fall of the Labor Share and the Rise of Superstar Firms, *Quarterly Journal of Economics*, 135: 645–709.

Barkai S. [2020], Declining Labor and Capital Shares, *Journal of Finance*, 75: 2421–2463.

Beaudreau B.C. [2020], *The Economics of Speed: Machine Speed as the Key Factor in Productivity*, Springer.

Beaudreau B.C., Lightfoot H.D. [2015], The Physical Limits to Economic Growth by R&D Funded Innovation, *Energy*: 84: 45–52.

Benzell S.G., Brynjolfsson E. [2019], *Digital Abundance and Scarce Genius: Implications for Wages, Interest Rates, and Growth*, Working Paper, MIT Initiative on the Digital Economy.

Benzell S.G., Kotlikoff L.J., LaGarda G., Sachs J.D. [2015], *Robots Are Us: Some Economics of Human Replacement*, Working Paper No. 20941, NBER.

Berg A., Buffie E.F., Zanna L.-F. [2018], Should We Fear the Robot Revolution? (The Correct Answer is Yes), *Journal of Monetary Economics*, 97: 117–148.

Bloom N., Jones Ch.I., Van Reenen J., Webb M. [2020], Are Ideas Getting Harder to Find?, *American Economic Review*, 110(4): 1104–1144.

Bostrom N. [2014], *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press.

Brynjolfsson E., McAfee A. [2014], *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W.W. Norton & Co.

Brynjolfsson E., Rock D., Syverson Ch. [2019], Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics, in: Agrawal A., Gans J., Goldfarb A. (eds.), *The Economics of Artificial Intelligence: An Agenda*: 23–57, University of Chicago Press.

Caselli F., Manning A. [2019], Robot Arithmetic: New Technology and Wages, *American Economic Review: Insights*, 1: 1–12.

Davidson T. [2021], *Could Advanced AI Drive Explosive Economic Growth?*, Report, Open Philanthropy.

Davies A. et al. [2021], Advancing Mathematics by Guiding Human Intuition with AI, *Nature*, 600: 70–74.

DeCanio S.J., [2016], Robots and Humans – Complements or Substitutes?, *Journal of Macroeconomics*: 49: 280–291.

Eloundou T., Manning S., Mishkin P., Rock D. [2023], *GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*, 2303.10130. arXiv.

Frey C.B., Osborne M. [2017], The Future of Employment: How Susceptible Are Jobs to Computerisation?, *Technological Forecasting and Social Change*: 114: 254–280.

Galor O. [2005], From Stagnation to Growth: Unified Growth Theory, in: Aghion P., Durlauf S.N. (eds.), *Handbook of Economic Growth*: 171–293, North-Holland.

Gordon R.J. [2016], *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*, Princeton University Press.

Graetz G., Michaels G. [2018], Robots at Work, *Review of Economics and Statistics*, 100: 753–768.

Groth Ch., Koch K.-J., Steger Th. [2010], When Economic Growth is Less Than Exponential, *Economic Theory*, 44: 213–242.

Growiec J. [2022a], *Accelerating Economic Growth: Lessons from 200 000 Years of Technological Progress and Human Development*, Springer.

Growiec J. [2022b], Automation, Partial and Full, *Macroeconomic Dynamics*, 26: 1731–1755.

Growiec J. [2022c], R&D Capital: An Engine of Growth, *Economics Letters*, 217: 110703.

Growiec J. [2023], What Will Drive Global Economic Growth in the Digital Age?, *Studies in Nonlinear Dynamics and Econometrics*, 27: 335–354.

Growiec J., Jabłońska J., Parteka A. [2023], *Hardware and Software over the Course of Long-Run Growth: Theory and Evidence*, KAE Working Paper 2023–091, SGH Warsaw School of Economics.

Growiec J., McAdam P., Mućk J. [2023], *R&D Capital and the Idea Production Functio*n, Research WP 23–05, Federal Reserve Bank of Kansas City.

Hanson R., [2001], *Economic Growth Given Machine Intelligenc*e, Working Paper, University of California, Berkeley.

Hemous D., Olsen M. [2018], *The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality*, Working Paper, University of Zurich.

Hidalgo C. [2015], *Why Information Grows: The Evolution of Order from Atoms to Economies*, New York: Basic Books.

Hilbert M., López P. [2011], The World's Technological Capacity to Store, Communicate, and Compute Information, *Science*, 332: 60–65.

Jones B.F. [2009], The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?, *Review of Economic Studies*, 76: 283–317.

Jones C.I. [1995], R&D-Based Models of Economic Growth, *Journal of Political Economy*, 103: 759–784.

Jones C.I. [2005], Growth and Ideas, in: Aghion P., Durlauf S.N. (ed.), *Handbook of Economic Growth*, North-Holland.

Jones C.I., Kim J. [2018], A Schumpeterian Model of Top Income Inequality, *Journal of Political Economy*, 126: 1785–1826.

Jones L.E., Manuelli R.E. [1990], A Convex Model of Equilibrium Growth: Theory and Policy Implications, *Journal of Political Economy*, 98: 1008–1038.

Korinek A. [2023], *Language Models and Cognitive Automation for Economic Research*, NBER Working Paper No. 30957.

Korinek A., Juelfs M. [2022], *Preparing for the (Non-Existent?) Future of Work*, Working Paper, University of Virginia.

OpenAI [2023], *GPT-4 Technical Report*, 2303.08774. arXiv.

Ord T. [2020], *The Precipice: Existential Risk and the Future of Humanity*, Hachette.

Philippon Th. [2022], *Additive Growth*, Working Paper No. 29950, National Bureau of Economic Research.

Piketty Th. [2014], *Capital in the Twenty-First Century*, Harvard University Press.

Romer P.M. [1986], Increasing Returns and Long-Run Growth, *Journal of Political Economy*, 94(5): 1002–1037.

Romer P.M. [1990], Endogenous Technological Change, *Journal of Political Economy*, 98: S71–102.

Rosling H. [2018], *Factfulness: Ten Reasons We're Wrong About the World – and Why Things Are Better Than You Think*, Sceptre.

Sachs J.D., Benzell S.G., LaGarda G. [2015], *Robots: Curse or Blessing? A Basic Framework*, Working Paper No. 21091, NBER.

Schwab K. [2016], *The Fourth Industrial Revolution*, World Economic Forum.

Solow R.M. [1956], A Contribution to the Theory of Economic Growth, *Quarterly Journal of Economics*, 70: 65–94.

Trammell Ph. [2021], *Existential Risk and Exogenous Growth*, Global Priorities Institute, University of Oxford.

Trammell Ph., Korinek A. [2021], *Economic Growth under Transformative AI: A Guide to the Vast Range of Possibilities for Output Growth, Wages, and the Labor Share*, Global Priorities Institute WP No. 8, Oxford University.

Wei J. et al. [2022], Emergent Abilities of Large Language Models, *Transactions on Machine Learning Research*, 8.

Yudkowsky E. [2013], *Intelligence Explosion Microeconomics*, Technical report No. 1, Machine Intelligence Research Institute.

Zeira J. [1998], Workers, Machines, and Economic Growth, *Quarterly Journal of Economics*, 113: 1091–1117.