August 2024

Working Paper

Hardware and software: A new perspective on the past and future of economic growth

Jakub Growiec, Julia Jabłońska, and Aleksandra Parteka

This working paper is available online at: https://www.brookings.edu/center/center-on-regulation-and-markets/

Disclosure

This research was supported by grants from the National Science Centre of Poland, OPUS 14 No. 2017/27/B/HS4/00189 and OPUS 19 No. 2020/37/B/HS4/01302. Other than the aforementioned, the authors did not receive financial support from any firm or person for this article or from any firm or person with a financial or political interest in this article. The authors are not currently an officer, director, or board member of any organization with a financial or political interest in this article.

Hardware and Software: A New Perspective on the Past and Future of Economic Growth*

Jakub Growiec[†] Julia Jabłońska[‡] Aleksandra Parteka[§]

August 9, 2024

Abstract. We propose a "hardware–software" framework that offers a new perspective on the mechanisms of long-run economic growth. Based on first principles, it assumes that output is generated through purposefully initiated physical action. Production needs energy and information, provided by respective factors: *hardware* ("brawn"), including physical labor, physical capital and compute, and *software* ("brains"), encompassing human cognitive work and digital software, in particular artificial intelligence (AI). Hardware and software are essential and complementary in production, whereas their constituent components are mutually substitutable. The framework generalizes the neoclassical model with capital and labor, models with capital–skill complementarity and skill-biased technical change, and selected unified growth theories. We provide an empirical quantification of hardware and software in the U.S., 1968–2019, documenting a rising share of physical capital in *hardware* (mechanization) and digital software in *software* (automation); as a whole software has been growing systematically faster than hardware. Accumulation of human capital and digital software were the key contributors to U.S. economic growth. Looking into the future through the lens of the hardware–software framework, we expect full automation of production by transformative AI and an order-of-magnitude acceleration of economic growth.

Keywords: production function, complementarity, mechanization, automation, artificial intelligence, transformative AI.
JEL codes: O30, O40, O41.

*Financial support from the Polish National Science Center (Narodowe Centrum Nauki) under grants OPUS 14 No. 2017/27/B/HS4/00189 (theory) and OPUS 19 No. 2020/37/B/HS4/01302 (evidence) is gratefully acknowledged. Thanks to Michał Gradzewicz, Anton Korinek and Tamay Besiroglu for helpful suggestions. All errors are our responsibility.

[†]SGH Warsaw School of Economics, Poland. ORCiD: 0000-0003-2222-1691.
[‡]SGH Warsaw School of Economics, Poland. ORCiD: 0000-0003-1280-3014.
[§]Gdańsk University of Technology, Poland. ORCiD: 0000-0003-1149-6614.

I'm a physicist. We rank things by two parameters: energy and information.

Michio Kaku

1 Introduction

In any conceivable technological process, output is generated through physical action requiring energy. It is a local reduction of entropy, and as such it does not occur by chance but is purposefully initiated. In other words, producing output requires both some physical *action* and some *code*, a set of instructions describing and purposefully initiating the action. Therefore, at the highest level of aggregation the two essential and complementary factors of production are physical *hardware* ("brawn"), performing the action, and disembodied *software* ("brains"), providing information on what should be done and how.

This basic observation has profound consequences. It underscores that the fundamental complementarity between factors of production, derived from first principles of physics, is cross cutting the conventional divide between capital and labor. From the physical perspective, it matters whether it's energy or information, not if it's human or machine (Figure 1). For any task at hand, physical capital and human physical labor are fundamentally substitutable inputs, contributing to hardware: they are both means of performing physical action. Analogously, human cognitive work and digital software are also substitutes, making up the software factor: they are alternative sources of instructions for the performed action. It is hardware and software, not capital and labor, that are fundamentally essential and mutually complementary.

Based on this observation the current paper develops a new macroeconomic framework for modelling aggregate production and long-run economic growth. We then demonstrate how it squares with historical data for the U.S. in 1968–2019 and what predictions it provides for the future.

Unfortunately, in data the fundamental distinction between hardware and software is obscured by the fact that the human body has double duty: it contains both muscles that perform physical action and a brain that stores and processes information. When performing any task, we make use of both energy and information, with varying intensity. The same can be said for modern digital devices, such as computers, smartphones and robots, which also feature both hardware and software. Prior to digital computers, though, all instructions were coming from the human brain, making "software" synonymous with human cognitive work. Therefore, while providing an overarching theoretical frame capable of

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

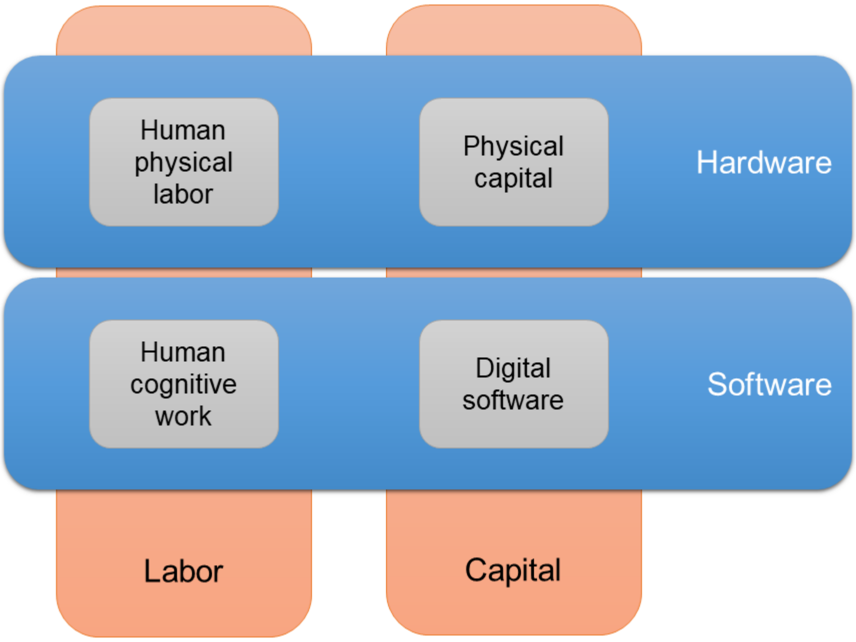

guiding the narrative across all human history (Growiec, 2022*a*), the advantages of the hardware–software framework are most clearly seen in the case of the currently unveiling digital era where information processing, communication and storage are increasingly detached from the human brain.

The hardware–software framework has a number of distinctive advantages. From the economic modelling perspective, it is a convenient tool for discussing global long-run growth processes because, while rooted in first principles from physics, it nests the following conventional models as special cases:

(i) the standard model of an industrial economy which uses capital and labor (Solow, 1956) and respects Kaldor (1961) facts (this case is obtained by assuming that all physical action is performed by machines, whereas all information processing is done by people),

(ii) a model of capital–skill complementarity and skill-biased technical change (this case is obtained assuming that all information processing is done by people),

(iii) a unified growth theory addressing the period of Industrial Revolution (following the arrival of machines with an external source of energy, able to perform physical action),

(iv) a theory of inception and further development of the digital era (following the arrival of programmable digital hardware – compute – and digital software, able to process information).

In the policy perspective, the hardware–software framework can inform the debate on the future of global economic growth – whether we should expect secular stagnation (Jones, 2002; Gordon, 2016; Gomułka, 2023), balanced growth with limited automation – "race against the machine" (Acemoglu and Restrepo, 2018), accelerated growth with disruptive automation (Brynjolfsson and McAfee, 2014; Brynjolfsson, Rock and Syverson, 2019; Trammell and Korinek, 2021) or technological singularity (Kurzweil, 2005; Bostrom, 2014; Roodman, 2020). In the baseline scenario, the hardware–software framework predicts an acceleration of economic growth, driven by broad-based AI-driven automation, culminating in the emergence of transformative AI (Growiec, 2023*b*). The *software* factor is expected to gradually decouple from human cognitive work and become proportional to compute instead because digital software can be virtually costlessly copied and thus can easily scale up to the level of available compute. Under constant returns to scale and in the absence of further technological revolutions[1], this would gradually reduce the role of technical change augmenting human cognitive work and eventually generate long-run endogenous growth by hardware (specifically, compute) accumulation alone. In the limit, all production will be automated.

Having laid out the theory, we quantify its predictions empirically, using U.S. data for 1968–2019. The empirical approach of the current study is to construct a time series for *hardware*, consisting of human physical labor and physical capital, and *software*, consisting of human cognitive work and digital software. To this end, we decompose labor into its physical (manual) and cognitive components, as well as isolate the hardware and software parts of capital investment. Our calculations assume an exogenous rate of technological progress which, in line with the theoretical setup, takes place in the domain of information and therefore is *software-augmenting*.

We find a rising share of physical capital in *hardware* (mechanization) and digital software in *software* (automation) throughout the period 1968–2019. On top of that, as a whole *software* has been growing systematically faster than *hardware*. Using a nested constant elasticity of substitution (CES) production function specification, we also perform a growth accounting exercise which suggests that the leading contributor to GDP growth in the U.S. has been the accumulation of human capital, followed by the accumulation of digital software.

Finally, we present the predictions of the hardware–software framework for the future. We lay out the range of possible scenarios and argue why the scenario

[1] Given the observed pace of growth in computing power and AI capabilities, further technological revolutions are actually quite likely, though.

allowing for full automation through transformative AI is our baseline. We conclude by discussing the implications of transformative AI for growth and factor remuneration.

This paper is related to at least five strands of literature. First, the literature on production function specification and estimation, in particular with capital–skill complementarity, unbalanced growth, as well as investment-specific and skill-biased technical change.[2] Second, the literature on accounting for the accumulation of information and communication technologies (ICT) and their broad growth-enhancing role as a general purpose technology.[3] Third, studies focusing on automation and its impacts on productivity, employment, wages and factor shares.[4] Fourth, the literature on macroeconomic implications of development of AI and autonomous robots.[5] Last but not least, the voluminous literature on R&D-based endogenous growth.[6]

The remainder of the paper is structured as follows. Section 2 provides motivation for the current study. Section 3 defines the factors of production in the hardware–software framework and discusses the conceptual underpinnings of the aggregate production function. Section 4 provides the empirical evidence. Section 5 provides model-based predictions for the future. Section 6 concludes.

---

[2]Including among others Gordon (1990); Jorgenson (1995); Greenwood, Hercowitz and Krusell (1997); Hercowitz (1998); Kumar and Russell (2002); Koop, Osiewalski and Steel (1999, 2000); Krusell et al. (2000); Henderson and Russell (2005); Caselli and Coleman (2006); Klump, McAdam and Willman (2007, 2012); Mućk (2017); McAdam and Willman (2018).

[3]Including among others Bresnahan and Trajtenberg (1995); Timmer and van Ark (2005); Jorgenson (2005); Brynjolfsson and McAfee (2014); Gordon (2016); Brynjolfsson, Rock and Syverson (2019); Aum, Lee and Shin (2018); Jones and Tonetti (2020); Farboodi and Veldkamp (2019); Nordhaus (2021).

[4]Including among others Zeira (1998); Acemoglu and Autor (2011); Autor and Dorn (2013); Graetz and Michaels (2018); Acemoglu and Restrepo (2018, 2019*a*,*b*); Andrews, Criscuolo and Gal (2016); Arntz, Gregory and Zierahn (2016); Frey and Osborne (2017); Barkai (2020); Autor et al. (2020); Jones and Kim (2018); Hemous and Olsen (2018); Benzell and Brynjolfsson (2019).

[5]Including among others Yudkowsky (2013); Graetz and Michaels (2018); Sachs, Benzell and LaGarda (2015); Benzell et al. (2015); DeCanio (2016); Acemoglu and Restrepo (2018); Aghion, Jones and Jones (2019); Berg, Buffie and Zanna (2018); Korinek and Stiglitz (2019); Trammell and Korinek (2021); Davidson (2021); Sevilla et al. (2022); Eloundou et al. (2023); Besiroglu, Emery-Xu and Thompson (2022); Erdil and Besiroglu (2023); Korinek and Suh (2024).

[6]Including among others Romer (1990); Jones and Manuelli (1990); Aghion and Howitt (1992); Jones (1995); Acemoglu (2003); Ha and Howitt (2007); Madsen (2008); Bloom et al. (2020); Kruse-Andersen (2023).

2 Motivation

2.1 New Trends of the Digital Era

Since the 1980s pre-existing long-run trends in economic development like Kaldor's "stylized facts" (Kaldor, 1961) and the seemingly eternal constancy of "great ratios" (Klein and Kosobud, 1961) have been overturned, and new ones emerged (Jones and Romer, 2010). Among the new tendencies, during the last 40 years we have been witnessing declining labor shares (Arpaia, Pérez and Pichelmann, 2009; Elsby, Hobijn and Sahin, 2013; Karabarbounis and Neiman, 2014), increasing profit shares (Barkai, 2020), increasing markups and market power (De Loecker, Eeckhout and Unger, 2020; De Loecker and Eeckhout, 2018; Diez, Leigh and Tambunlertchai, 2018), increasing market concentration (Autor et al., 2020) and increasing within-country income inequality (Piketty, 2014; Piketty and Zucman, 2014; Milanović, 2016). All this was accompanied by a tendency of skill polarization, gradual elimination of routine jobs (Acemoglu and Autor, 2011; Autor and Dorn, 2013), an increasing variety of jobs becoming susceptible to automation (Frey and Osborne, 2017; Arntz, Gregory and Zierahn, 2016; Eloundou et al., 2023), and a slowdown in total factor productivity growth (Jones, 2002; Gordon, 2016).

These emerging new tendencies can be understood as implications of the Digital Revolution which is transforming the world before our eyes in a comparably profound way to what the Industrial Revolution was doing two centuries ago. The computer age – differently from what Robert Solow observed back in 1987 – is now seen everywhere, even in the productivity statistics. Since the 1980s personal computers have been permeating firms and households, and digitization gained momentum in the 2000s with the spread of the Internet, and later – smartphones and AI. Quantitatively, since the 1980s "general-purpose computing capacity grew at an annual rate of 58%. The world's capacity for bidirectional telecommunication grew at 28% per year, closely followed by the increase in globally stored information (23%)" (Hilbert and López, 2011, p. 60). The costs of a standard computation have been declining by 53% per year on average from 1940 to 2014 (Nordhaus, 2021). Hence, growth in the digital sphere is now an order of magnitude faster than growth in the global capital stock and GDP: data volume, processing power and bandwidth double every 2–3 years, whereas global GDP doubles every 20–30 years (Growiec, 2022*a*). The processing, storage, and communication of information has decoupled from the cognitive capacities of the human brain; "less than one percent of information was in digital format in the mid-1980s, growing to more than 99% today" (Gillings, Hilbert and

Kemp, 2016, p. 183). New evidence also suggests that since the 1980s the efficiency of computer algorithms has been improving at a pace that is of the same order of magnitude as accumulation of digital hardware (Grace, 2013; Hernandez and Brown, 2020). Corroborating this finding, in the recent decade we have witnessed a surge in AI breakthroughs based on the methodology of *deep neural networks* (Tegmark, 2017): high-quality language interpretation, understanding, rephrasing, summarizing and producing human-like text (OpenAI's GPT-4, OpenAI, 2023, Anthropic's Claude, Google Gemini), generative visual art (SORA, DALL-E 3, Stable Diffusion), increasingly autonomous vehicles, self-taught superhuman performance at chess and Go (AlphaZero, Silver et al., 2018), or accurate prediction of protein structures (AlphaFold, Jumper et al., 2021). In recent years algorithmic progress in areas of AI such as language models, computer vision, and simulated gameplay, has been almost matching the growth in associated compute. Since 2012 language models "require 2x less compute roughly every eight months" (Ho et al., 2024, p. 5), whereas the compute used to train these models doubles roughly every six months (Sevilla et al., 2022). We are also observing that ever since Bill Gates first topped the list of World's Billionaires in 1995, many of the biggest fortunes are made in the computer software business.

Reconciling the implications of the Digital Revolution with established historical evidence in a single unified framework remains a challenge for growth theory: (i) conventional (textbook) economic growth models are rooted entirely in the industrial era (e.g. Barro and Sala-i-Martin, 2003; Jones, 2005*a*; Acemoglu, 2009); (ii) *unified* growth theories (e.g., Galor and Weil, 2000; Galor, 2005, 2011) are capable of successfully dissecting the mechanisms of transition from a relatively stagnant agricultural to a fast growing industrial economy during the Industrial Revolution, but not the transition from an industrial to a digital economy; (iii) models of economic growth with automation and AI (e.g., Acemoglu and Restrepo, 2018; Benzell et al., 2015; Berg, Buffie and Zanna, 2018; Aghion, Jones and Jones, 2019; Korinek and Stiglitz, 2019; Korinek and Suh, 2024) address the latter transition but not the former.

2.2 Mechanization, Automation and AI

The hardware–software framework involves a sharp conceptual distinction between mechanization and automation:

- *Mechanization* of production consists in replacing human physical labor with machines within hardware. Large-scale mechanization is observed since the Industrial Revolution (≈1800 CE onwards). Mechanization applies to physical actions but not the instructions defining them.

- *Automation* of production consists in replacing human cognitive work with digital software within software. But for early forerunners, automation is observed since the Digital Revolution (≈1980 CE onwards) when information technologies first came into use as general purpose technologies (Bresnahan and Trajtenberg, 1995). Automation pertains to cases where a task, previously involving human thought and decisions, is autonomously carried out by machines without any human intervention.

The distinction between mechanization and automation is instrumental in addressing questions like "will humans go the way of horses?" (Brynjolfsson and McAfee, 2014), which is supposed to mean whether human work will be eventually fully replaced by machines. The answer is: as far as physical labor is concerned, we have long gone the way of horses; for cognitive tasks (for which horses are of no use) this has not been the case yet, but it may happen in the future in the scenario of full (AI-driven) automation of production. By the same token, it is false comfort to say that the history of the Industrial Revolution teaches us that when jobs are destroyed, new ones are bound to emerge: that period only tells us that when human physical labor is mechanized, additional workers will be demanded in cognitive occupations, but it tells us nothing about cognitive occupations being automated.

The hardware–software framework also helps disentangle the concepts of automation and *artificial intelligence (AI)*. AI algorithms are a special type of software that has the ability to improve its performance based on experience and data. This happens even under a static architecture of AI algorithms, though it is conceivable that in the future AI may also modify its own architecture while heading towards self-improvement. In principle automation does not need AI, and indeed has historically begun prior to AI. However AI can strongly accelerate automation by substituting human cognitive work in non-routine tasks (Brynjolfsson, Rock and Syverson, 2019; Eloundou et al., 2023). According to Agrawal, Gans and Goldfarb (2017), while computers drastically lowered the costs of computing (arithmetic), AI drastically lowers the costs of *prediction*.

All in all, AI algorithms provide drastic improvements in the applicability, efficiency, and versatility of software, but do not constitute a qualitative change in its function as means of providing instructions to programmable hardware. Hence, the framework does not envisage a separate "AI revolution", and rather sees AI development as a massive boost to the Digital Revolution which already began with the early computer hardware and software. In our view, AI is to the digital era what the development of electricity and internal combustion engines was to the industrial era: a second wave of key breakthroughs, forcefully accelerating the impact of the initial revolutionary technological ideas on the economy and society, but not a separate technological revolution (Gordon, 2016).

3 The Hardware–Software Framework

We postulate that output is generated through (i) purposefully initiated (ii) physical action. Based on this premise we posit that at the highest level of aggregation any production function should feature *hardware* X, performing the physical action using energy, and *software* S, providing the instructions using information. This leads to a general form of a production function:

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

where $\mathcal{F} : \mathbb{R}^2_+ \rightarrow \mathbb{R}_+$ 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. The degree of their complementarity is an open question; the plausible range spans from perfect complementarity (Leontief form) if just one method of producing output exists, to imperfect complementarity if producers are allowed to choose their preferred technology from a technology menu (Jones, 2005*b*; Growiec, 2013, 2018). Intuitively, a little substitutability is likely because the same outcome can sometimes be generated with more resources (larger X) but less efficient code (smaller S), or vice versa, but the fundamental complementarity should nevertheless prevail. One natural way to instantiate this assumption is to take a CES specification with an elasticity of substitution $\sigma \in (0, 1)$, cf. Klump, McAdam and Willman (2007, 2012). The particular CES form of the $\mathcal{F}$ function is however not necessary for the results.[7]

The specification (1) abstracts from raw materials, energy resources and data sets which are being used up in the production process. It works as if we assumed that they were given for free and in infinite supply, or at least that they were sufficiently cheap and abundant that they would never become a bottleneck (think, e.g., of the supply of solar energy). Relaxing this simplifying assumption is left for further research.

3.1 Factors of Production

Hardware X includes physical actions performed by both humans and machines. Hence, it encompasses both human physical labor L and the services of physical capital K, where the former variable excludes any know-how or skill of the

[7] For example, Growiec and Mućk (2020) propose a more flexible parametric framework that also allows the modeler to control whether the factors are gross substitutes or gross complements.

worker, and the latter excludes any code that could be stored and executed on the machine.

Software S, in turn, encompasses all useful instructions which stem from the available information, in particular the practical implementation of state-of-the-art technologies. Hence, it includes the skills and technological knowledge employed in human cognitive work, H, as well as digital software Ψ providing instructions to be performed by the associated compute.[8] Services of digital software Ψ may be provided by static programs such as operating systems, spreadsheets or word processors, but also AI algorithms, distinguished by their ability learn from data as well as potentially self-improve and self-replicate. It is assumed that there are no physical obstacles precluding digital software from performing any cognitive task available to a human (Yudkowsky, 2013; Dennett, 2017).

Within hardware, agents of physical action are substitutable. The extreme case of perfect substitutability reflects the idea that whatever it is that performs a given task, if the set of actions is the same then the outcome should be the same, too. The same logic applies to software: regardless of whether a set of instructions comes from a human brain or a digital information processing unit, if the actual information content of instructions is the same, then the outcome should be the same, too. Therefore, at the level of sufficiently disaggregated tasks, all forms of software should also be considered perfectly substitutable.

However, this intuitive property will not always smoothly aggregate to the macro level. To see this, it helps to view the specification (1) as a reduced form of a richer framework where hardware and software are used in performing heterogeneous tasks, and the overall supply of hardware and software is computed by aggregating over these tasks (Acemoglu and Restrepo, 2018, 2019*a*,*b*; Growiec, 2022*b*). In such a scenario imperfect substitutability between human and machine contributions to factors of production may ensue from heterogeneity and mutual complementarity of the tasks. A particularly important caveat in this regard is that the baseline hardware–software framework excludes *essential non-automatable* cognitive tasks and sub-tasks – which cannot be circumvented and for which human cognitive work is necessary. For example, if a cognitive task consists of two consecutive steps, the first of which can be performed by a computer algorithm but the latter only by a human, then digital software and human cognitive work will turn out complementary at the level of the whole task even if

[8]Compute, i.e., programmable hardware, consists typically of computers, robots, and other devices embodying digital chips. In principle, it does not have to be silicon-based. In fact the first pieces of non-biological programmable hardware were mechanical devices such as the Jacquard loom using punchcards, first invented in 1804.

they are perfectly substitutable within the two sub-tasks. This apparent complementarity disappears, however, once the task becomes fully automatable.[9]

In line with this discussion, we write the general form of a production function as:[10]

$$Output = \mathcal{F}(X, S) = \mathcal{F}(L + K, H + \Psi). \tag{2}$$

Each of the four factors of production L, K, H, Ψ has unique properties (Table 1).

- *Human physical labor* L is rivalrous and given in fixed supply per worker, $L = \zeta N$ where $\zeta \in [0, \bar{\zeta}]$ denotes the supply of physical labor per worker in a unit of time, expressed in physical capital units, and N is the total number of workers.

- *Physical capital* K is rivalrous but can be unboundedly accumulated per capita. Physical capital K consists of non-programmable machines and compute. The share of compute in total physical capital is denoted by χ (so that $\chi \in [0, 1]$).

- *Human cognitive work* H consists of three components: technological knowledge A, the average skill level h, and the number of workers N, as in $H = AhN$. Technological knowledge A, or the size of the "repository of codes" is non-rivalrous (Romer, 1986, 1990) and accumulable.[11] Per-capita skill levels h are rivalrous and bounded above, theoretically by the optimal code for performing a given task, but in practice by a much lower number $\bar{h} > 0$ due to the human inability to rewire our brains in order to perform cognitive tasks more efficiently (Yudkowsky, 2013) as well as more down-to-earth reasons like human mortality and decreasing returns to education.

[9]Note that in the established task-based automation literature (Zeira, 1998; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018; Aghion, Jones and Jones, 2019) the default situation is that tasks can be only partially automated, whereas in the hardware–software framework in principle tasks can be automated fully. Growiec (2022*b*) demonstrates that a shift from partial to full automatability of complex tasks is disruptive for the economy – the contribution of human cognitive work switches from essential and scarce to inessential and replaceable – and argues that in the future we may see more and more tasks fully automated with the advancement of AI.

[10]At the cost of less transparent notation, one can generalize equation (2) to accomodate imperfect substitutability between people and machines in both hardware and software (Growiec, 2023*b*): $Output = \mathcal{F}(G_1(L, K), G_2(H, \Psi))$, with gross substitutability of factors within G_1 and G_2. A particularly tractable case to consider is the one where $\mathcal{F}, G_1$ and G_2 are CES. Furthermore, the partial automatability scenario – where some essential tasks will never be automated – can be accomodated by assuming gross complementarity between human and machine inputs in G_2 (Growiec, 2022*b*). This is the specification we use in the empirical Section 4 of our study, covering a historical period during which the potential for automation was limited.

[11]Depending on the institutional setup (e.g., intellectual property rights), technological knowledge A may be characterized by varying levels of excludability.

- *Digital software* Ψ also consists of three components: technological knowledge A, algorithmic skill level ψ which captures the degree to which digital software is able to perform the tasks collected in A, and the stock of compute χK on which the software is run, as in $\Psi = A\psi\chi K$. Technological knowledge A is the same as above.[12] The algorithmic skill level ψ is assumed to be bounded above by the optimal code for performing a given task (i.e., perfect accuracy), though there may be in fact a much lower upper bound $\bar{\psi}$ (Hanson and Yudkowsky, 2013).[13] Because software can be virtually costlessly copied, it is assumed that it can scale up to the level of all available compute χK.[14]

Table 1: Factors of Production and R&D

Hardware X	Human physical labor	$L = \zeta N$
	Non-programmable physical capital	$(1-\chi)K$
	Compute	χK
Software S	Human cognitive work	$H = AhN$
	Digital software†	$\Psi = A\psi\chi K$

Note: † includes AI algorithms.

3.2 Technological Progress

Following Romer (1986, 1990), the hardware–software framework envisages technological progress (growth in A) as an expansion of the "repository of codes", i.e., as the development of new, better instructions allowing to produce higher output with a given amount of hardware. Whether these new instructions take the form of new abstract ideas, scientific theories, systematically catalogued facts, codes specifying certain actions, or blueprints of physical items, they are all *information* and not actual *objects* or *actions*, and it is precisely this informational character that makes technologies non-rivalrous and a source of increasing returns to scale (Romer, 1990). What is novel here in comparison to Paul Romer's seminal contributions, though, is that these instructions can be applied to the tasks at hand by

[12] If in reality the sets of codes available to humans and digital software are different, the discrepancy between the measures of both sets can be captured by the ratio ψ/h.

[13] Depending on the institutional setup (e.g., proprietary code vs. open source), the algorithmic skill level ψ may be characterized by varying levels of excludability.

[14] Which implies that, in its basic form, the framework abstracts from economic and legal constraints on the diffusion of digital software, such as the protection of intellectual property rights.

both humans and machines.[15]

The informational character and non-rivalry of technological ideas naturally classify them in the domain of software, or "brains".[16] Consequently, all technological progress is modeled as *software-augmenting*.

The principle behind this sharp postulate is that the purpose of hardware is to perform physical action and *work* in the physical (mechanical) sense cannot be better or worse, there can only be more or less of it. Thus, construction of a machine able to, for example, transport a bigger load in the same amount of time and using the same amount of fuel, or to perform more digital computations per second using the same amount of energy, translates into *accumulation*, not *augmentation* of capital K. In turn, better targeted physical action achieved thanks to, say, a more precise tool or a better organized production stream indicates not an improvement in hardware, but software – instructions initiating the physical actions.

To explain this point further, let us link back to the literature on the capital embodiment controversy (Solow, 1960; Greenwood, Hercowitz and Krusell, 1997; Hercowitz, 1998). In line with the majority of modern papers on capital accumulation, the hardware–software framework views capital as "putty–putty" rather than "putty–clay". After installation, new machines simply enter the stock of capital K and we no longer trace their vintage and history. Even though systematic improvements in the energy efficiency of material processes are observed over time – new generations of machines are typically more capable and/or less energy consuming, and particularly strong compounding improvements are observed in the case of compute – because all these new machines are rivalrous devices which require to be built first, quality improvements embodied in them

[15]In the growth literature, the technology level A is frequently interpreted as mass of product designs (in *increasing variety* models) or an aggregate quality index of produced goods (in *quality ladder* models), Barro and Sala-i-Martin (2003). Note also the difference between technological ideas and data: "Ideas and data are types of information. Following Romer (1990), an idea is a piece of information that is a set of instructions for making an economic good, which may include other ideas. Data denotes the remaining forms of information. It includes things like driving data, medical records, and location data that are not themselves instructions for making a good but that may still be useful in the production process, including in producing new ideas." (Jones and Tonetti, 2020, p. 2821) In contrast to Jones and Tonetti (2020) and Farboodi and Veldkamp (2019) the hardware–software framework does not include data as a factor in the production function. Instead, data, like energy, is tentatively assumed to be sufficiently cheap and abundant that it will never become a bottleneck in production.

[16]Moreover, the technology level A can also be considered to include a reduced form for a variety of mechanisms and institutions underlying the equilibrium allocation in a more general model. For example, markets, legal systems and culture also implement some algorithms for aggregating and processing information, and those algorithms can vary in efficiency.

are included in the volume (dollar value) of capital accumulation and do not constitute hardware-augmenting technical change.

The postulate of purely software-augmenting technical change is particularly important over the long run perspective because it rules out *unbounded* growth in energy efficiency of material processes, which in reality is physically impossible because it would contradict the laws of thermodynamics (see e.g. Beaudreau and Lightfoot, 2015). Specifically, the Landauer's principle provides a theoretical lower bound for energy consumption of computation (Erdil and Besiroglu, 2023). *Bounded* improvements in energy efficiency are possible, though. For example, a sizeable share of incremental productivity improvements achieved over the two centuries of the industrial era could be attributed to the increasing speed of physical actions, such as the speed of travel or velocity of moving parts in industrial machines. Simultaneous energy-saving innovation allowed us to achieve those gains with less-than-proportional increases in energy consumption. While these economies of speed have been already largely exhausted (Beaudreau, 2020), now we are observing a different type of acceleration, namely in the speed and energy efficiency of digital computation. Computing efficiency (the energy use of computers) has been halving every 1.5 years over the last 60 years (Roser, Ritchie and Mathieu, 2023). Looking into the future, there is room for more progress in this regard, but eventually – following at least the Landauer's principle – these new "economies of speed" must be limited as well. In fact, physical constraints may become practically binding much earlier; they have already slowed down the progress in miniaturization of digital chips (Waldrop, 2016).

This is important for macroeconomic dynamics because (physically impossible) unbounded hardware-augmenting technical change would have the potential to resolve the scarcity of compute in production over the long run limit, creating a self-reinforcing feedback loop that leads to explosive, super-exponential growth (Growiec, 2023*b*). Such explosive dynamic is obtained, for example, when unbounded technological progress is introduced into the AK model or its generalizations (Erdil and Besiroglu, 2023).

In contrast, there are no analogous bounds on the number of potentially useful ideas, technologies, algorithms or blueprints. Even though there are known limits on the asymptotic complexity of certain algorithm families, such that, e.g., the number of operations must grow at least as fast as some function of the input size, these limits do not constrain the space of ideas as a whole and – by extension – the scope for software-augmenting technical change.

3.3 The Aggregate Production Function

Since Solow (1956, 1957) it has become commonplace to take capital K and labor L as the key inputs of the aggregate production function. Furthermore, it has become a very frequent, if not default, practice to assume purely labor-augmenting (Harrod-neutral) technical change, as in $Y = F(K, AL)$. Of course, like any other aggregate production function specification, this is a simplification that disregards the fact that K and L are amalgamates of heterogeneous components (Temple, 2006). The key question is, though, whether this simplified form is sufficient for capturing the key macroeconomic facts. Evidence is mounting that it is no longer the case. From the literature[17] it is becoming clear that the capital–labor framework, while sufficient to model the classic Kaldor (1961) facts, fails at capturing the new phenomena specific to the digital era, present in macro data since the 1980s.

In contrast, the hardware–software production function (2) specifies the production factors in accordance with the physical divide between energy and information, "brawn" and "brains". Using the concepts introduced above, the aggregate production function F is formalized as:

$$Y = F(X, S) = F(\zeta N + K, A(hN + \psi\chi K)), \tag{3}$$

where Y is aggregate value added (or GDP). The function F is increasing and concave in both its arguments, and hardware X and software S are essential and complementary. The standard replication argument applied to this production function specification implies constant returns to scale with respect to the rivalrous factors X and $S/A = hN + \psi\chi K$. With respect to X, S/A and A, though, returns to scale are increasing (Romer, 1986, 1990).

From the laws of thermodynamics, implying in particular that performing physical action requires expediting energy, it is expected that an essential fraction of GDP must consist of material outputs, serving – at the very least – to sustain the hardware (including human bodies) and allow it to work (Georgescu-Roegen, 1971, 1975). This observation reinforces the assumption that hardware X must be essential in the production process.

Digital software can be deployed in production processes only if the technology allows for the existence of compute ($\chi > 0$). Once it is introduced, though, there is no upper bound for its capacity relative to the cognitive capacity of the human brain. It may even one day come to exhibit superhuman cognitive per-

[17] Such as Gordon (1990); Greenwood, Hercowitz and Krusell (1997); Krusell et al. (2000); Caselli and Coleman (2006); Klump, McAdam and Willman (2007); Jones and Romer (2010); McAdam and Willman (2018).

formance.[18] This is because (i) the human brain has fixed computational capacity whereas digital software (including AI) can use arbitrarily large amounts of compute, (ii) AI algorithms have the ability to learn from data and potentially self-improve their architecture. Nevertheless, even without superhuman AI performance all cognitive tasks are amenable to automation with sufficient computing power χK. The only pre-condition for this outcome is that in the full model the possibility of accumulating the requisite computing power is not precluded by, e.g., preferences or institutions.[19]

It is instructive to consider four special cases of the framework, representing four distinct conventional models.

Industrial economy producing with capital and labor. Under the assumption that all physical work is done by machines ($\zeta = 0$) and all cognitive work is done by humans ($\chi = 0$), the production function (3) reduces to the conventional capital–labor specification with purely labor-augmenting technical change, $Y = F(K, AhN)$. Capital and labor are then naturally gross complements, as suggested by bulk of the recent empirical literature (Klump, McAdam and Willman, 2007, 2012; Mućk, 2017).

Capital–skill complementarity and skill-biased technical change. Under the assumption that all cognitive work is done by humans ($\chi = 0$), the production function (3) reduces to the specification with capital-skill complementarity (Krusell et al., 2000; Caselli and Coleman, 2006; McAdam and Willman, 2018) and skill-biased (or more precisely, cognitive labor-augmenting) technical change, $Y = F(\zeta N + K, AhN)$. Gross complementarity between hardware and software implies that physical capital is complementary to cognitive ($\approx$ skilled) labor H but substitutable with physical ($\approx$ unskilled) labor L, in line with findings of this literature.

Industrial Revolution. The hardware–software framework represents the Industrial Revolution as an episode where physical capital begins to be accumulated after the initial restriction $K \approx 0$ is lifted.[20] As a result, human physical labor is

[18] See Chollet (2019) for an excellent review of definitions of *intelligence* (cognitive performance, cognitive capabilities, etc.) of non-human agents.

[19] However, in a more general model with complex, multi-step tasks, human cognitive work can become essential for generating output if at least one step of at least one essential task is not automatable (Growiec, 2022*b*). Essentiality implies that there is no way around this particular step and no possibility of substituting out the entire task.

[20] The initial restriction $K \approx 0$ can be understood as the absence of machines with their own energy source (e.g., engine), able to perform physical action without energy inputs from the human.

gradually replaced by machines within hardware in a process of *mechanization* of production.

Digital Revolution. The framework represents the Digital Revolution as an episode where digital software begins to be accumulated after the initial restriction $\chi = 0$ (and thus $\Psi = 0$) is lifted. As a result, human cognitive work is gradually replaced with machine code within software in a process of *automation* of production.

3.4 Production Function for Ideas

Consistently with the hardware–software framework, research and development (R&D) processes are also viewed as a function of hardware X and software S. Hardware includes R&D capital alongside human physical labor (Growiec, 2022*c*; Growiec, McAdam and Mućk, 2023). Software encompasses all the ideas supplied by scientists and technical personnel, as well as code encapsulated in digital software. Formally the idea production function obeys the general equation (2), specializing into:

$$\dot{A} = \Phi(X, S) = \Phi(\zeta N + K, A^{\phi}(hN + \psi\chi K)), \tag{4}$$

where $\dot{A}$ represents the flow of new technological ideas. It is assumed that the idea production function Φ is increasing and concave in both factors, X and S.

Following the discussions in voluminous past literature (Jones, 1999; Ha and Howitt, 2007; Madsen, 2008; Bloom et al., 2020; Kruse-Andersen, 2023), we include software-augmenting knowledge spillovers in the production function for ideas, as represented by the parameter ϕ. The empirical magnitude of this parameter is subject to ongoing dispute. In theory, four distinct cases can be considered. First, if $\phi = 1$, then the stock of technological knowledge is included in the software term exactly like in the production function (3), so that no knowledge spillovers are present. With $\phi > 1$ there are positive knowledge spillovers, and with $\phi < 1$ – negative knowledge spillovers. Specifically, if $\phi < 0$ then the negative knowledge spillovers are so strong that they imply the emergence of "fishing out effects", due to which, *ceteris paribus*, R&D output declines with technological knowledge A (Growiec, 2023*a*).

3.5 Stages of Economic Development

Let us now trace how the hardware–software framework squares with the key properties of production processes across the human history (Growiec, 2022*a*). In this regard, it must be noted that the framework itself does not explain the causes of technological revolutions that push the economy to the next stage of development, other than speculating that in certain circumstances, given the relative supply of aggregate hardware vs. software, such a shift would be particularly demanded. However, the framework does predict the secular trends emerging after each technological revolution has exogenously occurred.

At this stage it is helpful to invoke the following asymptotic result:

$$F(1,\infty) = \lim_{y\to\infty} F(1,y) = a_X \in (0,+\infty). \tag{5}$$

Following from the assumptions of (i) constant returns to scale, and (ii) gross complementarity between hardware X and software S, the limit in (5) exists and is finite. One cannot achieve unbounded output growth unless both hardware and software grow unboundedly as well.

Stage 1. Pre-industrial production. In a pre-industrial economy, output was produced primarily in farming. People could only access the energy transformed in natural processes, such as photosynthesis and human metabolism. Without machines powered by external energy sources, there was no significant accumulation of productive capital K. Output was produced with a technology that used only human physical labor for performing the physical actions and required also the services of land, a vital but essentially fixed[21] factor of agricultural production. There was also no digital software Ψ. Setting a constant $K = \tilde{K}$, representing land, and $\chi = 0$ in equation (3) yields the following simple formula:

$$Y = F(X,S) = F(\zeta N + \tilde{K}, AhN) \approx N \cdot F(\zeta, Ah), \tag{6}$$

where the last approximation follows from the assumption that $\tilde{K}$ is fixed and small relative to ζN. Hence, due to gross complementarity of hardware and software, pre-industrial output per worker was bounded above ($Y/N \leq \zeta a_X$). The key growth bottleneck was the insurmountable scarcity of hardware (land and human physical labor), impossible to get around even in the hypothetical case $A \to \infty$.

Stage 2. Industrial production. Following the Industrial Revolution (≈1800 CE onwards) human physical labor was gradually replaced with steam-, oil-, and electricity-powered machines in production. The stock of physical capital per worker K/N began to grow exponentially. Productive physical actions were,

[21] By making this assumption we concentrate on a mature agricultural economy and exclude the periods of transition from hunting and gathering to sedentary agriculture or conquests of new agricultural land.

however, still dependent solely on the instructions produced through human cognitive work; there was no compute and no digital software yet. As hardware was accumulated faster than software, the latter eventually became relatively scarce, at which point demand for human cognitive skills began to grow, setting up a secular upward trend in wages (Galor, 2005). Setting $\chi = 0$ in (3) yields:

$$Y = F(X, S) = F(\zeta N + K, AhN). \tag{7}$$

The hypothetical limit of full mechanization and skill satiation, but with no automation, $K \to \infty$ and $h \to \bar{h}$, where $\bar{h}$ is the upper limit of human capital (skill) accumulation, implies $Y = F(K, A\bar{h}N)$. Hence for a mature industrial economy we obtain the standard balanced growth path result (Uzawa, 1961; Acemoglu, 2003). Under gross complementarity of capital and labor (really: hardware and software) and "labor-augmenting" (really: software-augmenting) technical change, the industrial economy tends to a balanced growth path where capital per worker K/N and output per worker Y/N grow at the same rate as technological knowledge A. Technological progress augmenting human cognitive work, generated through R&D, is the unique source of long-run growth (Romer, 1990). However, in contrast to the pre-industrial period, exponential growth is now sustained by the proportional expansion in complementary physical capital.

Stage 3. Digital production. Following the Digital Revolution (≈1980 CE onwards) we are observing gradual automation of production ($\chi > 0$). Human cognitive work which scales with the working population N is gradually replaced with digital software which scales with compute χK that grows faster. Consequently, software-augmenting technical change no longer affects only the efficiency of human cognitive work, but also to an increasing degree the capacities of digital software. As automation progresses, skill-biased technical change morphs into routine-biased technical change (Acemoglu and Autor, 2011; Autor and Dorn, 2013). This is the world in which we live now.

At a later stage of the digital era, however, digital software will likely consist largely of advanced general-purpose AI algorithms, allowing for multiple-fold increases in the algorithmic skill level ψ (Agrawal, Gans and Goldfarb, 2017; Berg, Buffie and Zanna, 2018) and thus fortifying the emerging upward trend in the contribution of the non-human component to software.

The limit of $K \to \infty$ and $\psi \to \bar{\psi}, \chi \to \bar{\chi}$ implies

$$Y = F(\zeta N + K, A(hN + \psi\chi K)) \approx K \cdot F(1, A\bar{\psi}\bar{\chi}), \tag{8}$$

where $\bar{\psi}$ is the upper limit of algorithmic skill accumulation and $\bar{\chi} \in (0, 1]$ is the limiting share of compute in all physical capital as $K \to \infty$. Full automation of production in the limit implies that the human contribution to output will gradually fall to zero.[22]

Equation (8) delivers an AK-type implication: in contrast to the industrial economy, long-run growth of the digital economy is driven not by technological progress but by the accumulation of hardware (specifically, compute) (Jones and Manuelli, 1990; Barro and Sala-i-Martin, 2003). If $A \to \infty$ then $Y/K \to a_X$. This striking result is driven by two forces: (i) that digital software expands proportionally with compute, and (ii) that hardware and software are gross complements, and thus in the long run the pace of accumulation of hardware – the scarce factor – determines the pace of economic growth. The constancy of the output growth rate over the long run follows in turn from the assumption of constant returns to scale in production, making F asymptotically linear in K (Jones, 2005*a*; Growiec, 2007).

Although asymptotically constant, the pace of hardware accumulation and output growth may be nevertheless stupefying, potentially with doubling times of the order of 2–3 years, which are currently observed for digital processing power, data volume and bandwidth (Hanson, 2001; Hilbert and López, 2011).[23]

Hypothetical stage 4. Post-digital production. Under high to full automation of production, compute χK will gradually become the bottleneck of economic growth, the key factor constraining its pace. This will increase the incentives to invest in R&D directed towards radical innovations holding the promise to eliminate this bottleneck. Such breakthrough technology would have to tap an entirely new source of energy, fundamentally increase energy efficiency, or otherwise massively improve unit productivity of compute.[24]

Formally, such an episode of "new mechanization" (or "new compute") may be modelled by introducing an additional component to the hardware amalgamate, as in:

$$X = \zeta N + K + \omega M, \tag{9}$$

where M denotes the new form of hardware, and $\omega \gg 1$ captures its unit pro-

[22]Full automation does not necessarily mean that human work will one day become *useless* for the economy (Harari, 2017). The decline in human productivity relative to machines will surely be reflected in sub-par growth in wages, but the extent of technological unemployment will eventually depend also on the elasticity of labor supply. See also Korinek and Juelfs (2022).

[23]Long-run growth in the alternative scenario where some essential tasks will never be automated is investigated in Growiec (2023*a*,*b*).

[24]Among the probable scenarios, one could envision the arrival of quantum computing (in which case the Google AI Quantum team has already achieved a major breakthrough, Arute et al., 2019), disruptive nanotechnology, massively improved solar power cells, fusion power, or perhaps something yet unimagined.

ductivity relative to K. This form of hardware must be programmable, so that AI could scale with M and avoid becoming a growth bottleneck itself.

Long-run implications include gradual replacement of K with M and a permanent acceleration in growth. Indeed, this additional acceleration in hardware X accumulation may eventually lead to a new growth regime "with a doubling time measured in days, not years" (Hanson, 2000, p. 1): technological singularity.

3.6 Factor Shares

The assumption of gross complementarity of hardware and software provides a clear-cut implication for factor shares: factor income will be disproportionately directed towards the scarce factor. The hardware–software framework delivers the following (empirically testable and intuitively explicable) predictions.

Stage 1. Pre-industrial production. In a mature pre-industrial economy able to achieve systematic technological progress (growth in A), increasing scarcity of human physical labor and agricultural land ($\zeta N + \tilde{K}$) relative to human cognitive work (AhN) implies that an ever-increasing portion of value added is directed to hardware at the expense of software. The counterfactual limit of $A \to \infty$ without an industrial revolution (with a fixed $K = \tilde{K}$) implies a zero software share of output as virtually all revenues are directed towards agricultural (physical) workers and owners of agricultural land.

Stage 2. Industrial production. The first stage of development of an industrial economy features gradual *mechanization* of production: physical capital accumulation systematically reduces the role of human physical labor in hardware. Given the substitutability between capital K and physical labor ζN, the physical labor share goes down whereas the capital share goes up – a trend which was most clearly seen in the early 19th century and was reflected in Marx's ideas regarding the exploitation of the working class.

However, as the pace of capital accumulation in a growing industrial economy outruns technical change (growth in A), this secular trend is accompanied also by an increasing output share accruing to software (i.e., human cognitive work) at the expense of hardware ($\zeta N + K$, gradually dominated by K). Hence, during the second stage of development of an industrial economy, human cognitive work becomes increasingly scarce and thus increasingly well remunerated, raising the returns to education and the skill premium, and setting up a secular upward trend in wages. Such trend was observed in developed countries from the late 19th and throughout the 20th century.[25] In the counterfactual limit of $A \to \infty$,

[25] As Galor and Moav (2006) put it, "The accumulation of physical capital in the early stages

$K \to \infty$ and $h \to \bar{h}$ without a digital revolution, the industrial economy tends to a balanced growth path, along which $Y = F(K, A\bar{h}N)$, the hardware (=capital) share stabilizes at some intermediate value $\bar{\pi}_X \in (0, 1)$, and the economy respects Kaldor's facts (Kaldor, 1961).

Stage 3. Digital production. The first stage of development of a digital economy features gradual *automation* of production: accumulation of digital software Ψ gradually reduces the role of human cognitive work H in software. Given the substitutability of these two factors, the cognitive labor share goes down whereas the digital software share goes up. (And if data and software rents are not separately accounted, also firms' profit shares and measured markups go up, as documented e.g. by Barkai (2020); De Loecker and Eeckhout (2018).) This is the world of today, where booming digital technologies fuel the "rise of the global 1%".

The hardware-software framework predicts a change in this secular trend in the future, though. It expects that due to ongoing technological progress in A, systematic improvements in algorithmic skill ψ and progressing automation, hardware (and more precisely, compute) will gradually become the bottleneck of global development, a key factor constraining the pace of further economic growth.[26] Consequently the revenues will be increasingly redirected from software towards compute, and the software share π_S will set on a secular downward trend. In the hypothetical limit of $K \to \infty, \chi \to \bar{\chi}, \psi \to \bar{\psi}$, assuming the absence of a next technological revolution, the hardware share will tend to unity. At that point in time, though, only a negligible fraction of total remuneration will be earned by human workers.

4 Empirical Evidence for the USA, 1968–2019

Mapping the theoretical concepts of L, K, H and Ψ to real-world data is a challenge. In the data there is no direct split of workers' time and remuneration between their physical labor and cognitive work; each worker in some proportion does both. Similarly, programmable devices also have double duty as means of performing physical action and as compute which stores and runs its code; measured capital investment and returns conflate both. It is not even clear in

of industrialization enhanced the importance of human capital in the production process and generated an incentive for the capitalists to support the provision of public education for the masses, planting the seeds for the demise of the existing class structure" (p. 85).

[26]This is a robust prediction which fails only if full automation is not possible (then human cognitive work remains the growth bottleneck forever) or if there is also hardware-augmenting technical change (which leads to super-exponential, explosive growth), cf. Growiec (2023*b*).

the accounting whether a certain investment expenditure helps accumulate programmable or non-programmable capital. Finally, if not for intellectual property rights digital software can be virtually costlessly copied to a multiplicity of devices, making it notoriously difficult to price and evaluate its marginal productivity.

In this section we provide a first attempt at quantifying hardware and software, using U.S. data for the years 1968–2019. We construct the relevant time series and plug them into a growth accounting exercise. We use (i) O*NET Content Model data, providing detailed information on work characteristics and equipment used in almost 1000 occupational groups; (ii) microdata from the CPI IPUMS (Flood et al., 2022) on hours worked by occupation in the U.S. from 1968 to 2019; and (iii) tables on U.S. investment in fixed assets by category from the U.S. Bureau of Economic Analysis; (iv) aggregate GDP, hours worked and labor share in the U.S.

4.1 Decomposing Labor: Manual vs. Cognitive Tasks

Our first step is to isolate the hardware and software component within labor (L and H, respectively). To this end we decompose work tasks in individual professions into manual and cognitive tasks using the method proposed by Autor, Levy and Murnane (2003); Acemoglu and Autor (2011). However, while these seminal papers and the subsequent task-based literature (e.g., Spitz-Oener, 2006; Autor and Handel, 2013; Autor, Dorn and Hanson, 2015; Lewandowski et al., 2022) focused on the split between routine and non-routine task categories, we identify the *manual* (physical) vs. *cognitive* content of jobs. We merge raw O*NET (v.25.3) files on Work Activities, Work Context, Abilities and Skills and identify manual tasks using a specific list of selected Work Activities and Work Context Importance scales.[27] For each occupation, we measure the share of manual work as the average importance of manual tasks, while the share of cognitive work is obtained as a residual. In so doing we follow O*NET procedure of standardization into 0-100 scores (because in raw data, each separate task descriptor in O*NET is associated with a different scale).

These shares are then matched with occupation-level employment data. The shares of individual occupations in overall hours worked in the U.S. economy

[27] Routine manual: 4.C.3.d.3 Pace determined by speed of equipment; 4.A.3.a.3 Controlling machines and processes; 4.C.2.d.1.i Spend time making repetitive motions. Non-routine manual, physical adaptability: 4.A.3.a.4 Operating vehicles, mechanized devices, or equipment; 4.C.2.d.1.g Spend time using hands to handle, control or feel objects, tools or controls; 1.A.2.a.2 Manual dexterity; 1.A.1.f.1 Spatial orientation.

Figure 2: Dynamics of the share of manual and cognitive work in the U.S. (1968=1)

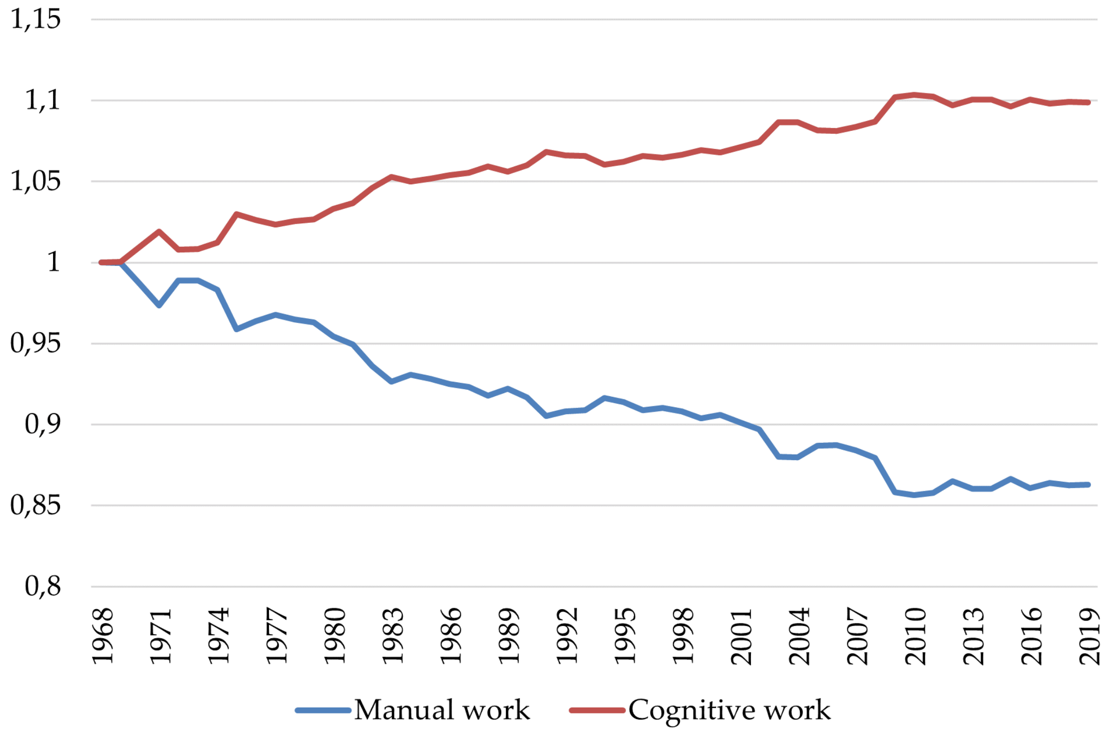

*Source: own computations based on O*NET and CPS IPUMS data.*

are extracted from the Current Population Survey (CPS) IPUMS database (Flood et al., 2022), containing microdata from the monthly U.S. labor force survey. The classification system used in the IPUMS database covers more than 450 occupations. We include observations of persons who were professionally active, had a specific occupation and disclosed the number of hours worked. To map the ~ 1000 occupations in O*NET with ~ 450 occupations in CPS IPUMS, we use the crosswalk O*NET-SOC 2019 to 2018 SOC from the O*NET Resource Centre. Upon aggregation, we obtain the split of labor between manual and cognitive work in the U.S. in the period 1968–2019 (Figure 2).[28]

Finally, we obtain our final time series corresponding to physical labor L (that enters hardware) and cognitive labor H (that enters software) by multiplying the shares of manual and cognitive work by total hours worked in the U.S. economy in a given year.

[28]Our method for splitting total hours worked into manual and cognitive tasks yields conservative estimates, with relatively little growth in the ratio of cognitive to manual work. This may be partly because, due to data limitations, we identify the manual vs. cognitive content of jobs at only one point in time, and all the measured temporal variation comes from changes in the occupational structure of employment.

Figure 3: Hourly wage in manual and cognitive work (USD, constant prices)

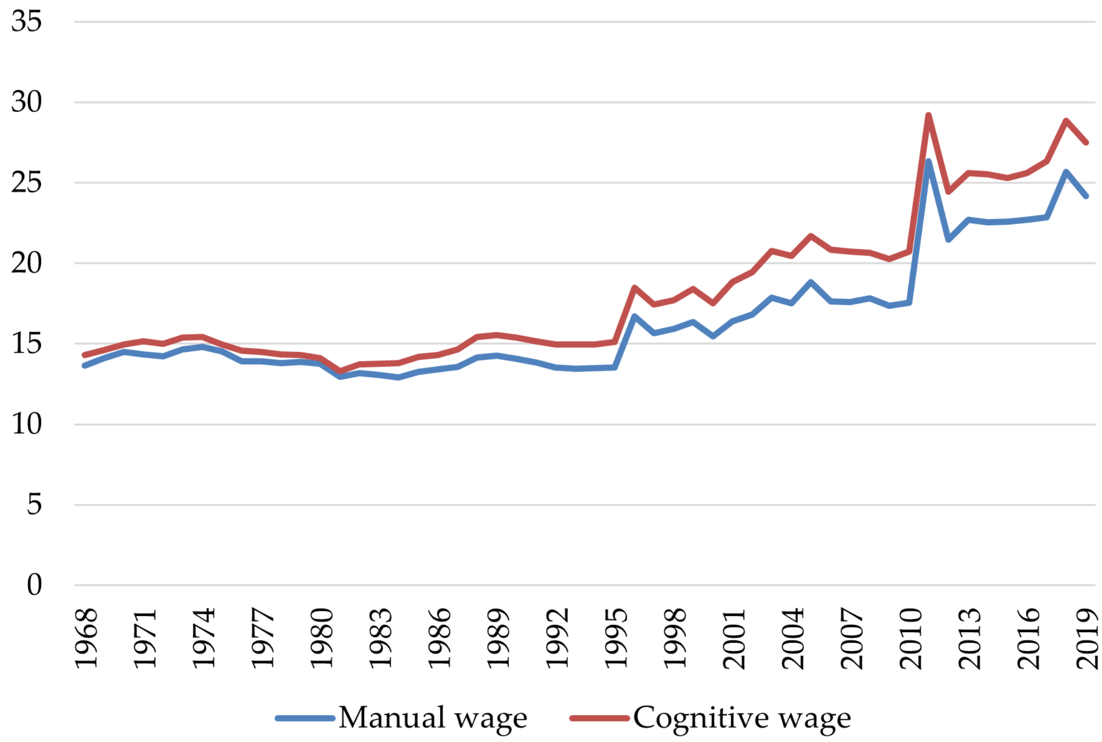

*Source: own computations based on O*NET and CPS IPUMS data.*

We also construct a time series of average real wages in manual and cognitive tasks, using CPS IPUMS data on remuneration by occupation as well as the aforementioned O*NET (v.25.3) dataset on Work Activities, Work Context, Abilities and Skills. The manual (respectively, cognitive) wage is calculated as a weighted average of hourly wages across occupations, using the total hours worked in performing manual (cognitive) tasks within each occupation as weights (Figure 3). The measured difference between hourly wages for performing physical labor and cognitive work is rather low – cognitive work pays about 10% more on average, with a slow increase in the premium over time – mirroring our conservative estimates of the split between manual vs. cognitive tasks within jobs.

4.2 Decomposing Capital: Physical Capital vs. Digital Software

The process of breaking down total capital into physical capital K (that enters hardware) and digital software Ψ is analogous. First, we take U.S. Bureau of Economic Analysis data which allows us to divide investment into structures, intellectual property products (IPPs) and 25 categories of equipment.[29] We assume that investment in structures contributes 100% to hardware, while investment in IPPs contributes 100% to software. The challenge, however, is to determine to what extent investments in specific types of equipment affect the hardware and software stock. In an attempt to solve this problem via proxy, we use O*NET data which provides information on the type of equipment used in the day-to-day work in various professions. We assume that the more manual the job is, the more hardware-intensive equipment the worker uses. In contrast, highly cognitive occupations are assumed to be more likely to use equipment containing mostly digital software.[30] We attribute to each type of equipment its specific proportion of hardware and software by merging the Tools Used by Occupation dataset from O*NET and the occupation-level manual-cognitive split discussed above. Using the O*NET-SOC 2019 codes we merge the Tools Used by Occupation dataset with BEA dataset on real investment by category.

As a result of these steps, we obtain a time series on real investments in physical capital (hardware) and digital software. Next, we use the standard perpetual inventory method to build up the stocks of physical capital (hardware) and digital software. We apply asset-specific depreciation rates based on Fraumeni (1997). These rates range from 0.026 per annum (structures) to 0.315 (computers and peripheral equipment).

For illustrative purposes, we also calculate a weighted average of asset-specific deflators, using the total hardware and software stock in each asset category as weights (Figure 4). The result is striking: hardware prices were rapidly growing throughout the entire time frame 1968–2019, whereas software prices were roughly constant. This finding is in line with the plentiful past evidence that the price of equipment relative to structures (which we count as 100% hardware) exhibits a secular downward trend (e.g., Greenwood, Hercowitz and Krusell, 1997; Gordon, 2016). Specifically, the relative prices of ICT equipment have been falling

[29] Private fixed assets; Computers and peripheral equipment; Communication equipment; Medical equipment and instruments; Nonmedical instruments; Photocopy and related equipment; Office and accounting equipment; Fabricated metal products; Engines and turbines; Metalworking machinery; Special industry machinery, n.e.c.; General industrial, including materials handling, equipment; Electrical transmission, distribution, and industrial apparatus; Trucks, buses, and truck trailers; Autos; Aircraft; Ships and boats; Railroad equipment; Furniture and fixtures; Agricultural machinery; Construction machinery; Mining and oilfield machinery; Service industry machinery; Electrical equipment, n.e.c.; Other nonresidential equipment; Residential equipment.

[30] This is a tentative assumption that calls for refinement in the future. Anecdotal evidence suggests that it is not always the case that cognitive tasks are performed with "smart" devices, and manual work – with simple tools. However, we do not have sufficient data to verify this.

Figure 4: Unit price of physical capital (hardware) and digital software (USD, constant prices)

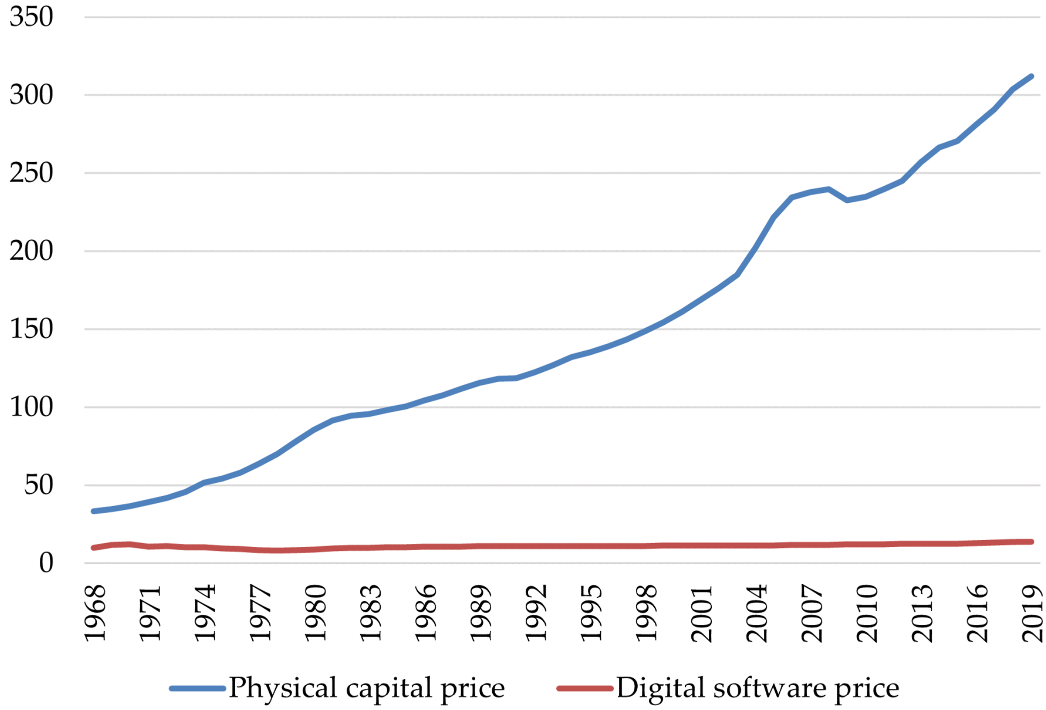

*Source: own computations based on O*NET and BEA data.*

most precipitously (Timmer and van Ark, 2005), and accordingly in the U.S. BEA data *computers and peripheral equipment* are the category that witnessed the most extreme price declines, not just in relative but also in absolute terms.

4.3 Constructed Time Series

Over the period 1968–2019 there was a clear parallel increase in the share of cognitive work in labor and of digital software in capital (Figure 5). Significant differences in software intensity have persisted between capital and labor, though: digital software accounted for about 48% of total capital in the U.S. in 2019 (in current prices), up by 5 pp. since 1968; at the same time human cognitive work (software) constituted about 64% of total labor input, up by 6 pp. since 1968.

At this point we must also posit a functional form for exogenous software-augmenting technical change $A(t)$ feeding into human cognitive work H and digital software Ψ. As this is the first attempt to quantify hardware and software, we opt to keep things as simple as possible. Therefore, we postulate exponential technological progress at a constant rate $g > 0$, i.e. $A(t) = e^{gt}$. In the baseline calibration (see below) we assume $g = 1.5\%$ per annum.

Figure 5: Share of human cognitive work in labor and digital software in capital (in %)

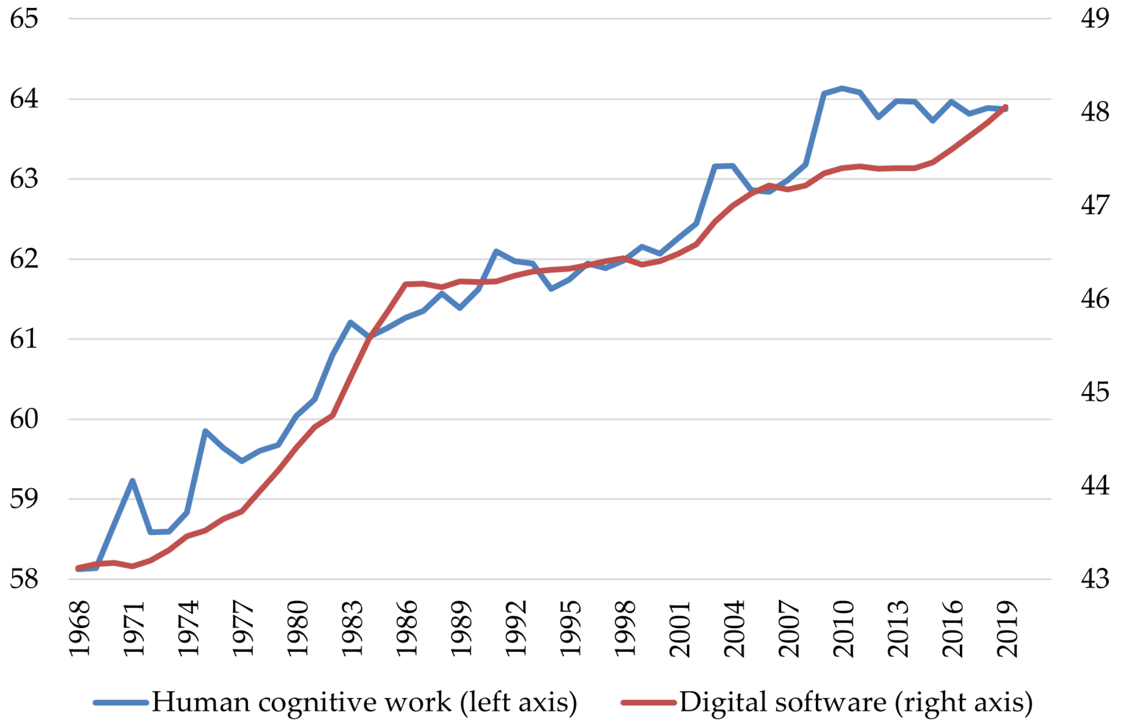

*Source: own computations based on O*NET, CPS IPUMS and BEA data.*

With this in hand we find that while the stocks of all four factors of production (L, K, H, Ψ) have been growing over time, in line with U.S. population growth and fixed asset formation, their growth rates were rather disparate (Figure 6). Human physical labor L was growing at 0.8% per annum on average, cognitive work H and real physical capital K – at 2.8%, and digital software Ψ – at 4.8% (compared to the average real GDP growth rate of 2.7% per annum[31]). Clearly, even without specifying the relative contribution of L vs. K in hardware and H vs. Ψ in software, we already see that as a whole software has been growing systematically faster than hardware.

4.4 Calibration of the Aggregate Production Function

We now combine all four factors of production in a modified version of aggregate production function (3). This aggregate production function will later be also used in a growth accounting exercise.

We use the nested normalized CES production function specification, with

[31] As calculated based on BEA data downloaded in June 2022. Revisions to the data may have occurred since then.

Figure 6: Dynamics of physical capital K, digital software Ψ, physical labor L and cognitive work H (1968=1)

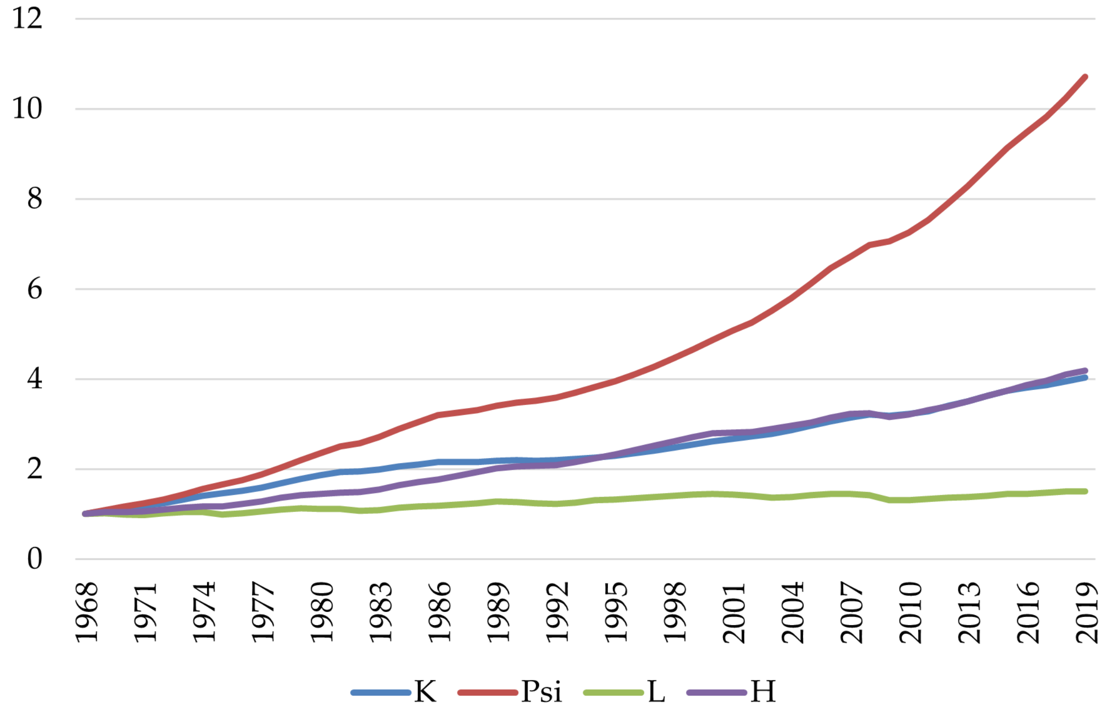

*Source: own computations based on O*NET, CPS IPUMS and BEA data.*

hardware and software being gross complements:

$$Y = Y_0 \left(\alpha \left(\frac{X}{X_0} \right)^{\theta} + (1-\alpha) \left(\frac{S}{S_0} \right)^{\theta} \right)^{\frac{1}{\theta}}, \qquad \theta < 0, \alpha \in (0,1). \tag{10}$$

In contrast to (3), we now also use normalized CES aggregates for hardware and software, thereby relaxing the assumption of perfect substitutability between people and machines within hardware and within software. We do so because in reality there is a multiplicity of tasks to be performed, both in terms of physical action and information processing; even if people and machines are perfectly substitutable in performing each task, the tasks themselves may be complementary and many tasks certainly have not been fully automatable in the considered time period (Growiec, 2022*b*). Hence, we write:

$$X = X_0 \left(\gamma \left(\frac{L}{L_0} \right)^{\mu} + (1-\gamma) \left(\frac{K}{K_0} \right)^{\mu} \right)^{\frac{1}{\mu}}, \qquad \mu \leq 1, \gamma \in (0,1), \tag{11}$$

$$S = S_0 \left(\beta \left(\frac{H}{H_0} \right)^{\omega} + (1-\beta) \left(\frac{\Psi}{\Psi_0} \right)^{\omega} \right)^{\frac{1}{\omega}}, \qquad \omega \leq 1, \beta \in (0,1). \tag{12}$$

In line with usual practices in the normalization literature (Klump, McAdam and

Willman, 2012), the normalization points with subscript 0 are taken as (geometric) sample means.

However, as the hardware–software framework is a new theoretical setup, there is no evidence in the literature on the values of distribution parameters α, β, γ, and elasticity parameters θ, μ, ω. We set them so as to roughly match the (i) the average GDP growth rate (2.7% in data), (ii) average labor share (0.61 in data), and (iii) the cognitive wage premium (in data, an average hour of cognitive work is worth $\sim 10\%$ more than an hour of manual work), while excluding parametrizations implying implausible variability in the labor share and the cognitive wage premium.

To compare the predictions of the hardware–software framework with observations on the U.S. labor share and cognitive wage premium, we need to derive their model-based counterparts. Postulating the normalized CES specification (10) and assuming that factors are priced at their respective marginal products (subject to a possible constant markup), we obtain the following factor shares:

$$\pi_X = \alpha \left(\frac{X}{X_0} \frac{Y_0}{Y} \right)^{\theta}, \qquad \pi_S = (1-\alpha) \left(\frac{S}{S_0} \frac{Y_0}{Y} \right)^{\theta}, \tag{13}$$

$$\pi_L = \gamma \left(\frac{L}{L_0} \frac{X_0}{X} \right)^{\mu}, \qquad \pi_H = \beta \left(\frac{H}{H_0} \frac{S_0}{S} \right)^{\omega}, \tag{14}$$

where π_X is the hardware share of output and π_S is the associated software share (due to constant returns to scale with respect to rivalrous inputs and software-augmenting technical change in (10), $\pi_X + \pi_S = 1$); π_L is the physical labor share of hardware, and π_H is the cognitive labor share of software.

Using this notation, the labor share and the cognitive wage premium are derived as follows:

$$\pi_{Labor} = \pi_X \pi_L + \pi_S \pi_H, \tag{15}$$

and

$$\frac{w_H}{w_L} = \frac{\pi_S \pi_H}{\pi_X \pi_L} \frac{L}{H}. \tag{16}$$

The calibrated baseline parameters approximately achieving the aforementioned objectives are listed in Table 2.

Table 2: Baseline parameterization of the nested CES production function

Output		Hardware		Software		Tech
α	θ	γ	μ	β	ω	g
0.44	-0.2	0.45	1	0.71	-1.74	0.015

In this parameterization, the elasticity of substitution between hardware and software is $\sigma_{X,S} = \frac{1}{1-\theta} = 0.83$ – somewhat above the usual estimate of the elasticity of substitution between aggregate capital and labor from the literature, $\sigma \approx 0.6$ (Klump, McAdam and Willman, 2012), but in the theoretically postulated domain of gross complementarity. In turn, physical capital and human physical labor are perfectly substitutable, whereas human cognitive work and digital software are gross complements, with an elasticity of substitution of $\sigma_{H,\Psi} = \frac{1}{1-\omega} = 0.36$. These parameter choices imply that at the point of normalization, cognitive work earns about 17% more than manual work, whereas the labor share is 0.60.

Figure 7: The stocks of hardware and software (1968=1, left axis) and the extent of mechanization and automation (1968=1, right axis) under the baseline calibration

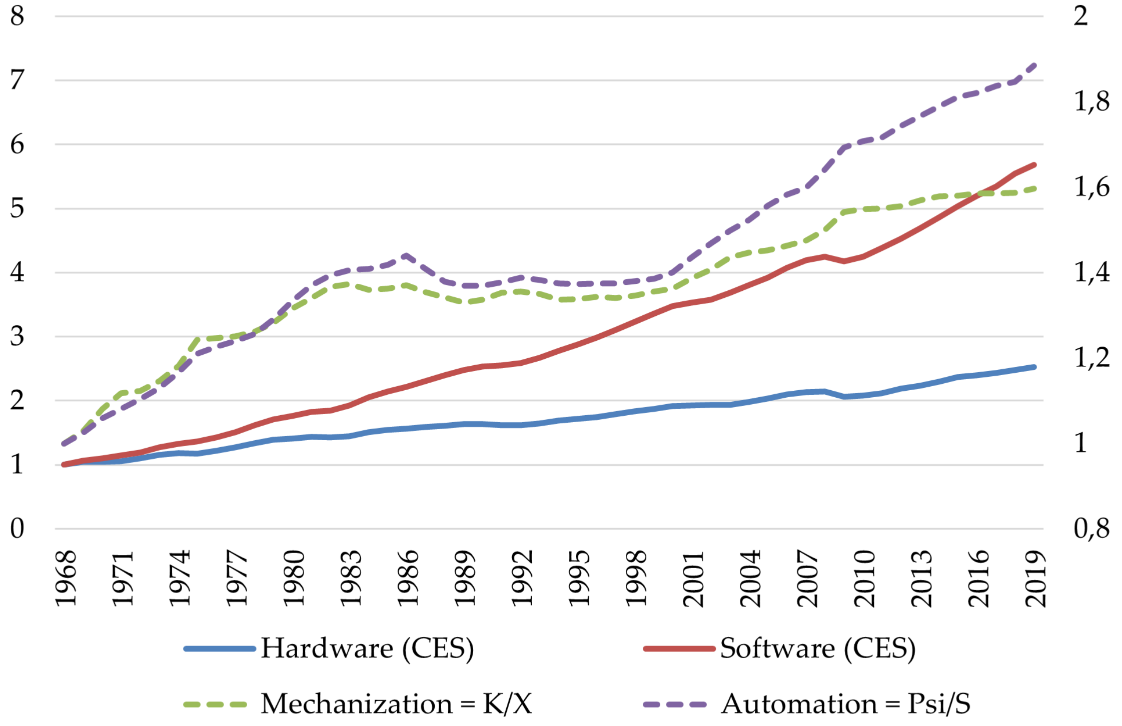

*Source: own computations based on O*NET, CPS IPUMS and BEA data.*

Our main results (Figure 7) suggest that growth in software (3.5% per annum) systematically outruns that of hardware (1.8% per annum). We also find that the extents of both mechanization (share of machines in hardware) and automation (share of machines in software) are upward trending – roughly in parallel until the early 2000s, but with a clear acceleration in automation afterwards. In total, over the considered time frame mechanization progressed by about 60% (or 0.9% per annum), and automation – by about 88% (1.3% per annum).

4.5 Growth Accounting

Log-differentiating equation (10) with respect to time, we obtain the following Solow-type decomposition of economic growth:

$$g_Y = \pi_X g_X + \pi_S g_S, \tag{17}$$

where $\pi_X = \frac{\partial Y}{\partial X}\frac{X}{Y}$ is the hardware share of output, and analogously $\pi_S = \frac{\partial Y}{\partial S}\frac{S}{Y}$ is the software share. Decomposing (10) further, we obtain:

$$g_Y = \pi_X \pi_L g_L + \pi_X \pi_K g_K + \pi_S \pi_H g_H + \pi_S \pi_\Psi g_\Psi, \tag{18}$$

where $\pi_L = 1 - \pi_K = \frac{\partial X}{\partial L}\frac{L}{X}$ is the human physical labor share within hardware, and $\pi_H = 1 - \pi_\Psi = \frac{\partial S}{\partial H}\frac{H}{S}$ is the human cognitive labor share within software.

Table 3: Contributions to annual GDP growth, 1968–2019

	GDP	K	Ψ	L	H	Residual
pp.	2.71	0.64	0.75	0.17	1.13	0.02
% of total		23.7%	27.9%	6.1%	41.7%	0.8%

*Source: own computations based on O*NET, CPS IPUMS and BEA data.*

Figure 8: GDP growth decomposition, 1968–2019

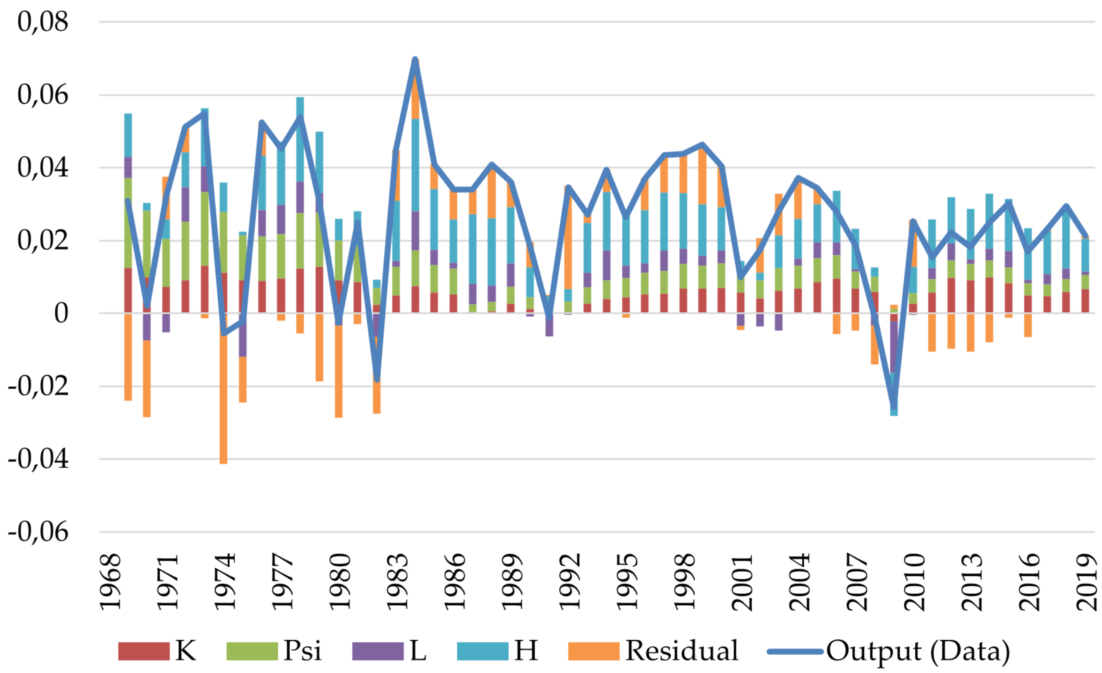

*Source: own computations based on O*NET, CPS IPUMS and BEA data.*

Under the baseline calibration (Table 2) we find that the key contributors to GDP growth in the U.S. in 1968–2019 were the accumulation of human capital H, followed by the accumulation of digital software Ψ and physical capital K (Table 3, Figure 8). Furthermore, while the contribution of human capital was roughly steady throughout the studied time period, the contribution of digital software was particularly strong in the 1970s and 1990s until mid-2000s. Our speculative hypothesis, to be verified as new data come along, is that the subsequent reduction in this factor's contribution to economic growth (from mid-2000s onwards) may constitute an interlude before the next upcoming wave of AI-driven automation in the coming years (Brynjolfsson, Rock and Syverson, 2019).

4.6 Robustness Checks

Alternative Characterization of Physical Capital vs. Digital Software. In the baseline scenario we decompose real investment into hardware and software investment based on O*NET data on equipment used across jobs. However, the identifying assumption of proportionality: the more manual the job is, the more hardware-intensive equipment the worker uses, need not hold exactly. Hence, as a robustness check we relax this assumption and instead concentrate on the equipment used only in n "most manual" and "most cognitive" jobs. Given about 1000 occupations in the database, we consider $n = 100, 200, 300$.

As shown in Figure 9, this modification reduces the implied stocks of digital software compared to the baseline, particularly strongly for lower values of n. This is because there are many more jobs in the database which are (almost) entirely manual than jobs that are (almost) entirely cognitive. However, except for the case $n = 100$ which is likely most noisy, the dynamics of hardware vs. software accumulation change only very slightly. Therefore, our results regarding the preferred parametrization of the aggregate production function as well as growth accounting are robust to these changes.

Figure 9 also shows that weighting jobs by hours worked in the decomposition of real investment into hardware and software slightly increases the estimated share of software. However, as the dynamics of hardware vs. software accumulation change minimally, our main results remain robust.

Evolution of Work Tasks Over Time. In the baseline scenario we identify the manual vs. cognitive content of jobs at only one point in time, and all the measured temporal variation in the shares of manual and cognitive work comes from changes in the occupational structure of employment. However, in fact the task content of jobs has been evolving over time. Specifically, as documented by Spitz-Oener

Figure 9: The share of digital software in capital under alternative characterizations (in %)

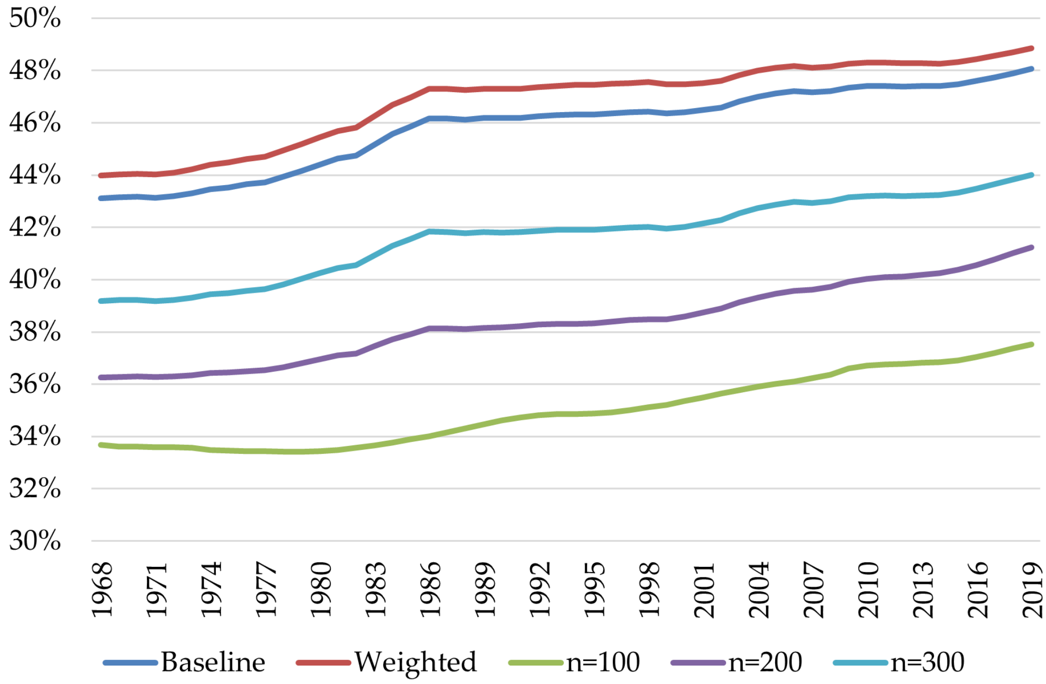

*Source: own computations based on O*NET, CPS IPUMS and BEA data.*

(2006) based on a unique dataset from West Germany that begins in the 1970s, over decades there has been a systematic decline in routine tasks and a systematic increase in nonroutine tasks within predetermined job titles. Notwithstanding, given that both routine and nonroutine tasks can be either manual or cognitive, the direction of impact of these trends on the implied shares of manual vs. cognitive content of jobs is *ex ante* unclear and requires further scrutiny.

Hence, as a robustness check we augment our calculation of percentages of manual and cognitive work performed in the U.S. in each year with Spitz-Oener (2006) data on systematic variation in job tasks, broken down into routine manual, routine cognitive, nonroutine manual, nonroutine cognitive, and nonroutine analytic tasks. We make the assumptions that (i) the data on task shares within jobs, collected at four points in time, can be linearly interpolated across all years and extrapolated beyond the time span of the dataset,[32] and (ii) the task content of jobs identified for West Germany fully applies also to the U.S. as well.

[32]The observed changes in shares of job tasks are monotone over time and the linear fit is quite good in-sample. However, theoretically one could expect these linear trends to break down out-of-sample, particularly when the predicted task shares approach 0% or 100%. As we do not include any "saturation" effects, the current results should be viewed as an upper bound for the extent of actual change in the task content of jobs.

As shown in Figure 10, the current modification radically changes our assessment of the dynamics of the software component of both capital and labor. The share of human cognitive labor now rises by as much as 29 pp. in the period 1968–2019 (compared to 6 pp. in the baseline), whereas the share of digital software in capital now rises by 34 pp. (rather than just 5 pp.). This means that empirically, the shift from routine towards nonroutine tasks within jobs has been accompanied by a systematic, sharp increase in cognitive tasks compared to manual ones. Concurrently, given our approach to identifying the share of digital software within capital, this shift is also reflected in systematic sharp increases in the digital software Ψ relative to physical capital K.

Figure 10: The share of human cognitive work in labor and digital software in capital (in %): baseline vs. time-variable tasks

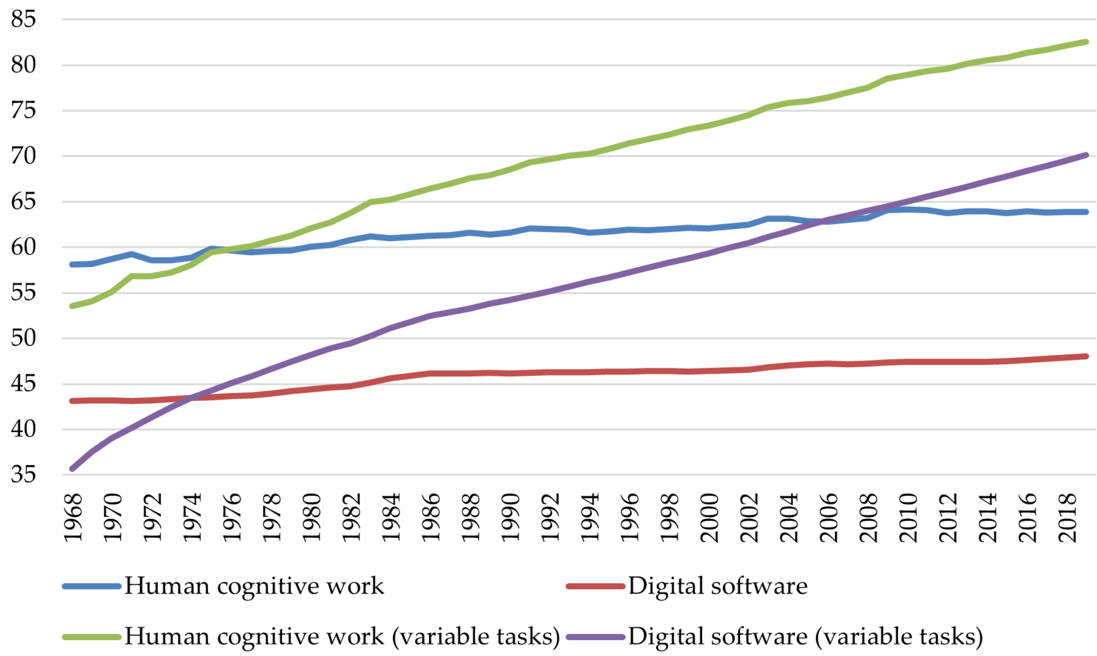

*Source: own computations based on O*NET, CPS IPUMS, BEA and Spitz-Oener (2006) data.*

Notwithstanding this radical change, the calibration of the nested CES production function remains remarkably robust. The elasticity of substitution between hardware and software remains at $\sigma_{X,S} = \frac{1}{1-\theta} = 0.83$. Physical capital and human physical labor are still assessed as perfectly substitutable, whereas human cognitive work and digital software remain gross complements, now with an elasticity of substitution of $\sigma_{H,\Psi} = \frac{1}{1-\omega} = 0.39$. The biggest difference compared to the baseline is observed with regard to the implied pace of software-augmenting technical change, which drops from 1.5% per annum in the baseline to 0.7% per annum now. This is because now a sizeable share of technical change is subsumed by the observed shifts in the task content of jobs and the associated increases in the software component of capital.

Like in the baseline scenario, we find that growth in software (now at 3.9% per annum) systematically outruns that of hardware (0.7% per annum). The extents of both mechanization (share of machines in hardware) and automation (share of machines in software) are upward trending – roughly in parallel until the early 2000s, but with a clear acceleration in automation afterwards. In total, over the considered time frame mechanization progressed by about 82% (or 1.2% per annum), and automation – by 124% (1.6% per annum).

Table 4: Contributions to annual GDP growth, 1968–2019

			Baseline			
	GDP	K	Ψ	L	H	Residual
pp.	2.71	0.64	0.75	0.17	1.13	0.02
% of total		23.7%	27.9%	6.1%	41.7%	0.8%
			Time variable tasks			
	GDP	K	Ψ	L	H	Residual
pp.	2.71	0.34	1.28	-0.11	1.14	0.06
% of total		12.4%	47.4%	-4.1%	42.4%	2.3%

*Source: own computations based on O*NET, CPS IPUMS, BEA and Spitz-Oener (2006) data.*

As far as growth accounting results are concerned (Table 4), we now see a substantially larger contribution of digital software Ψ and reduced contributions of physical capital K and human physical labor L. In fact, the contribution of human physical labor is now slightly negative because the total supply of physical labor is now assessed as gradually falling over time as jobs are becoming less and less manual in nature.

5 Predictions for the Future

The hardware–software framework constitutes a theoretical frame that can be used to discuss both the past and the future of the world economy. With regard to the future, it allows us to formulate a range of scenarios that would ultimately place the world economy on a spectrum: secular stagnation – balanced growth – accelerated growth – technological singularity. However, the exact prediction depends on few key assumptions whose validity is *ex ante* uncertain.

The key question is whether in the future human cognitive work and digital software will become gross substitutes, or if they will forever remain gross complements as they are today. This is the question of partial vs. full automation (Growiec, 2022*b*). Under full automation, human cognitive work could be fully replaced by digital software. That would unpin the economic growth rate from growth in the capacities of technologically augmented labor, and instead pin it to the growth rate of compute. Simultaneously, under full automation progress in AI capabilities and accumulation of compute would act to reduce, rather than increase the labor share of output as they currently do under partial automation.

The full automation scenario is therefore highly transformative. With today's knowledge it appears to rest on the possibility of artificial superintelligence, or transformative AI. Its implications will be discussed in a separate subsection.

5.1 Overview of Scenarios

Let us now review the possible scenarios in the order of increasing long-run growth rates.

- *Secular stagnation.* As suggested among others by Jones (2002); Gordon (2016), output growth may gradually slow down in the future, perhaps heading towards a steady state, driven by declining population growth and exhaustion of the pool of new technological ideas. However, in the hardware–software framework secular stagnation is possible only in a highly restrictive scenario, and even in that scenario growth continues forever, albeit at declining rates (i.e., growth is sub-exponential).
  - **Assumptions.** Full automation of production is impossible. Knowledge spillovers in R&D are negative ($\phi < 1$).
  - **Implications.** Because some essential production tasks cannot be automated, human cognitive work and digital software are gross complements. Therefore, regardless of the algorithmic skill level ψ and available compute χK, human cognitive work remains the bottleneck of economic growth. Growth is then driven solely by technological progress augmenting human cognitive work, and the rate of that progress is slowing down over time due to negative knowledge spillovers.
- *Balanced growth.* In a similar scenario, differing only in the magnitude of knowledge spillovers in R&D, the hardware–software framework predicts that balanced growth is maintained throughout the digital era. The long-run growth rate remains similar to the ones observed today (i.e., about

$2-3\%$ per annum, Piketty, 2014) or is mildly larger. This is a "race against the machine" scenario (Acemoglu and Restrepo, 2018): tasks are never fully automated, and labor-augmenting technological progress together with accumulation of R&D capital form a dual growth engine (Growiec, 2023*a*).

  - **Assumptions.** Full automation of production is impossible. Knowledge spillovers in R&D are positive or zero ($\phi \geq 1$).
  - **Implications.** Again, because some essential production tasks cannot be automated, human cognitive work and digital software are gross complements. Therefore, regardless of the algorithmic skill level ψ and available compute χK, human cognitive work remains the bottleneck of economic growth. Growth is driven by technological progress augmenting human cognitive work. The rate of that progress is sustained, in the absence of population growth, thanks to the ongoing accumulation of R&D capital.

- *Accelerated growth.* The predictions of the hardware–software framework change dramatically once full automation of production is allowed. As human cognitive work and digital software become gross substitutes, progress in the algorithmic skill level (ψ) – representing specifically AI capabilities – and the accumulation of compute χK lead to gradual replacement of people with machines within software. People are then going to be employed only as long as their services are cheaper than that of AI. This is the baseline case of the hardware–software framework discussed in Section 3.

  - **Assumptions.** Full automation is possible.
  - **Implications.** Human cognitive work and digital software eventually become gross substitutes. From then on, the accumulation of digital software (in particular, AI) resolves the scarcity of software in the aggregate production function. Eventually, hardware becomes relatively scarce because it is not technology-augmented. Economic growth is then driven by the accumulation of compute, the scarce factor complementary to the fast-growing capabilities of AI. In the long run, the rate of economic growth rate is equated with the growth rate of compute. Were compute to continue growing at the pace of Moore's Law, i.e., about $20-30\%$ per annum (Hilbert and López, 2011), compared to $2-3\%$ average growth in global output (Piketty, 2014), this would imply a growth acceleration by an order of magnitude.

- *Technological singularity.* Hypothetically, under full automation, there could be a further technological revolution in the domain of hardware, able to alleviate the mounting scarcity of compute. That would trigger another growth acceleration, potentially by another order of magnitude or more, resulting in technological singularity. The likelihood of this scenario relative to the former one depends on the extent of returns to cognitive reinvestment in AI, i.e., its potential for self-improvement (Yudkowsky, 2013), and the (now unknown) difficulty of achieving the next technological breakthrough in hardware.

  – **Assumptions.** Full automation is possible. Moreover, a new form of programmable hardware M arrives, gradually replacing existing compute χK.

  – **Implications.** AI switches from using compute χK to the new hardware M. In the long run, economic growth is proportional to growth in M. Over the transition, it is conceivable that AI may need to reprogram itself to be compatible with the new form of hardware. Even more hypothetically, if for some reason some essential tasks could not be reprogrammed, then compute χK would remain the bottleneck of economic growth, and the hypothesized new wave of growth acceleration would not occur.

5.2 Implications of Transformative AI

Under the baseline specification of the hardware–software framework, all essential tasks will eventually be automated, making human cognitive work replaceable, and thus paving the way for massive growth acceleration. Implicit in this scenario is that AI will one day not only exceed human capabilities at each narrow task, but also will be sufficiently versatile and adaptive to handle the aggregation and management of those tasks, strategic decision making, and R&D tasks leading to the creation of new technologies, new tasks, and new AI capabilities. In other words, it subsumes the future arrival of transformative AI (Trammell and Korinek, 2021; Korinek and Suh, 2024). Following a variety of philosophical, information-theoretic, anthropological and economic arguments (Yudkowsky, 2013; Bostrom, 2014; Tegmark, 2017; Growiec, 2022*a*) as well as the broadly shared belief among AI experts (Grace et al., 2024), we expect this scenario to be not only possible but actually quite likely – and hence we consider it our baseline.

The prospective arrival of transformative AI is associated with a number of fundamental questions, such as the AI alignment problem, existential risk from misaligned transformative AI, the evolution of income inequality in the absence of paid labor, or the value of human life in a world ruled by superhuman AI. Of course, the hardware–software framework is too simplified to inform any of these issues; however, it does offer systematic predictions on economic growth, factor shares, and transformation of the labor market in such a world. Namely the framework predicts that:

1. *Transformative AI will accelerate economic growth, likely by an order of magnitude.* Output growth, once unpinned from the growth rate of labor productivity, will eventually reach the growth rate of compute. Note that over the recent decades compute has been growing according to the Moore's Law, i.e., at $20-30\%$ per annum, doubling every 2-3 years (Hilbert and López, 2011).

2. *Human cognitive work will be substitutable with AI.* In a world with transformative AI, people will only find employment as long as they are price competitive against the AI. As demonstrated by Growiec (2022*b*), under perfect substitutability and perfect market competition, wages are then expected to stay constant despite rapid economic growth. Under imperfect substitutability, wages could grow, but still at a systematically lower rate than output. This opens up the possibility of technological unemployment; however, the market outcome will depend also on the elasticity of labor supply. It is conceivable that some people may be willing to work for any wage, especially if their livelihood will be secured through other means such as transfers or returns to proprietary compute.

3. *The labor income share will drop precipitously toward zero.* In a world with transformative AI, technological progress and accumulation of compute will be pushing the labor income share down toward zero as more and more jobs are automated. Once the majority of output is contributed by AI and the associated compute, wages will cease to be the key distributive device. Other devices will have to be sought, such as centralized redistribution of returns to compute (the scarce factor of production), or distributed ownership of compute.

6 Conclusion

In this paper, we have put forward the hardware–software framework – a new conceptual framework of production and long-run growth, based on first principles and emphasizing the role of energy and information. Nevertheless, it remains closely linked with the existing economic literature. It nests four conventional macro models as special cases and can be used to inform the debate on the future of global economic growth in the 21st century.

As an empirical application of the theory, we have constructed time series of physical capital K, digital software Ψ, physical labor L and cognitive work H for the U.S. in 1968–2019. We have then plugged these series into a growth accounting exercise. Our results suggest that the key contributor to GDP growth in the U.S. in 1968–2019 was the accumulation of human capital, followed by the accumulation of digital software. This is consistent with the interpretation (Growiec, 2022*a*) that we are still at an early stage of the digital era, and more profound economic transformations should be expected as AI-driven automation gains steam and more and more production processes are fully automated, thereby reducing the contribution of human cognitive work towards zero (Brynjolfsson, Rock and Syverson, 2019; Korinek and Stiglitz, 2019; Growiec, 2022*b*; Korinek and Juelfs, 2022; Eloundou et al., 2023).

Our results can be extended in a number of directions. First, one can build formal macroeconomic models based on the hardware–software framework, with a variety of applications. For example, Growiec (2023*b*) applied the hardware–software framework to build scenarios for the future and address the question: what will drive global economic growth in the digital age? Second, using certain identifying assumptions one can construct time series for hardware and software stretching further back in time, thus quantifying the role of these fundamental factors of production over the very long run, including for example the period of Industrial Revolution. This is needed to ascertain usefulness of the framework as a building block in a unified growth theory (Kremer, 1993; Galor, 2005, 2011). Third, one can add more detail to the model, such as heterogeneous tasks with varying extents of automatability (Growiec, 2022*b*). This would improve the fit of the model to the data and make it better suited to producing quantitative predictions of economic growth at later stages of the digital era.

References

Acemoglu, Daron. 2003. "Labor-and Capital-Augmenting Technical Change." *Journal of the European Economic Association*, 1: 1–37.

Acemoglu, Daron. 2009. *Introduction to Modern Economic Growth.* Princeton University Press.

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*. Vol. 4, , ed. O. Ashenfelter and D. Card, Chapter 12, 1043–1171. Elsevier.

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment." *American Economic Review*, 108: 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. 2019*a*. "Artificial Intelligence, Automation, and Work." In *The Economics of Artificial Intelligence: An Agenda*. , ed. Ajay Agrawal, Joshua S. Gans and Avi Goldfarb. University of Chicago Press.

Acemoglu, Daron, and Pascual Restrepo. 2019*b*. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33: 3–30.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica*, 60: 323–351.

Aghion, Philippe, Ben F. Jones, and Charles I. Jones. 2019. "Artificial Intelligence and Economic Growth." In *The Economics of Artificial Intelligence: An Agenda*. , ed. Ajay Agrawal, Joshua Gans and Avi Goldfarb, 237–282. University of Chicago Press.

Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2017. "What to Expect From Artificial Intelligence." *MIT Sloan Management Review*, Spring 2017.

Andrews, Dan, Chiara Criscuolo, and Peter N. Gal. 2016. "The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective." OECD Working Party No. 1 on Macroeconomic and Structural Policy Analysis.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2016. "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis." OECD Publishing, Paris OECD Social, Employment and Migration Working Paper No 189.

Arpaia, Alfonso, Esther Pérez, and Karl Pichelmann. 2009. "Understanding Labour Income Share Dynamics in Europe." European Commission, European Economy Paper No. 379 Working paper.

Arute, Frank, Kunal Arya, Ryan Babbush, et al. 2019. "Quantum Supremacy Using a Programmable Superconducting Processor." *Nature*, 574: 505–510.

Aum, Sangmin, Sang Yoon Lee, and Yongseok Shin. 2018. "Computerizing Industries and Routinizing Jobs: Explaining Trends in Aggregate Productivity." *Journal of Monetary Economics*, 97: 1–21.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–97.

Autor, David H., and Michael J. Handel. 2013. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages." *Journal of Labor Economics*, 31: S59–S96.

Autor, David H., David Dorn, and Gordon H. Hanson. 2015. "Untangling Trade and Technology: Evidence from Local Labour Markets." *Economic Journal*, 125: 621–646.

Autor, David H., David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics*, 135: 645–709.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics*, 118(4): 1279–1333.

Barkai, Simcha. 2020. "Declining Labor and Capital Shares." *Journal of Finance*, 75: 2421–2463.

Barro, Robert J., and Xavier X. Sala-i-Martin. 2003. *Economic Growth.* MIT Press.

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

Beaudreau, Bernard C., and H. Douglas Lightfoot. 2015. "The Physical Limits to Economic Growth by R&D Funded Innovation." *Energy*, 84: 45–52.

Benzell, Seth G., and Erik Brynjolfsson. 2019. "Digital Abundance and Scarce Genius: Implications for Wages, Interest Rates, and Growth." National Bureau of Economic Research NBER Working Paper No. 25585.

Benzell, Seth G., Laurence J. Kotlikoff, Guillermo LaGarda, and Jeffrey D. Sachs. 2015. "Robots Are Us: Some Economics of Human Replacement." NBER Working Paper No. 20941.

Berg, Andrew, Edward F. Buffie, and Luis-Felipe Zanna. 2018. "Should We Fear the Robot Revolution? (The Correct Answer is Yes)." *Journal of Monetary Economics*, 97: 117–148.

Besiroglu, Tamay, Nicholas Emery-Xu, and Neil Thompson. 2022. "Economic Impacts of AI-augmented R&D." arXiv 2212.08198.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review*, 110(4): 1104–1144.

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies.* Oxford University Press.

Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies: Engines of Growth?" *Journal of Econometrics*, 65: 83–108.

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

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2019. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In *The Economics of Artificial Intelligence: An Agenda*. , ed. Ajay Agrawal, Joshua Gans and Avi Goldfarb, 23–57. University of Chicago Press.

Caselli, Francesco, and Wilbur J. Coleman. 2006. "The World Technology Frontier." *American Economic Review*, 96: 499–522.

Chollet, Francois. 2019. "On the Measure of Intelligence." Google, Inc. Working paper.

Davidson, Tom. 2021. "Could Advanced AI Drive Explosive Economic Growth?" Open Philanthropy Report.

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

De Loecker, Jan, and Jan Eeckhout. 2018. "Global Market Power." National Bureau of Economic Research Working Paper 24768.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics*, 135: 561–644.

Dennett, Daniel. 2017. *From Bacteria to Bach and Back: The Evolution of Minds.* W. W. Norton & Company.

Diez, Federico J., Daniel Leigh, and Suchanan Tambunlertchai. 2018. "Global Market Power and Its Macroeconomic Implications." International Monetary Fund Working Paper No. 18/137.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." arXiv 2303.10130.

Elsby, Michael, Bart Hobijn, and Ayseful Sahin. 2013. "The Decline of the U.S. Labor Share." *Brookings Papers on Economic Activity*, 47(2): 1–63.

Erdil, Ege, and Tamay Besiroglu. 2023. "Explosive Growth From AI Automation: A Review of the Arguments." arXiv 2309.11690.

Farboodi, Maryam, and Laura Veldkamp. 2019. "A Growth Model of the Data Economy." MIT Sloan Working paper.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry. 2022. "Integrated Public Use Microdata Series." Current Population Survey: Version 10.0.

Fraumeni, Barbara M. 1997. "The Measurement of Depreciation in the U.S. National Income and Product Accounts." *Survey of Current Business, Bureau of Economic Analysis*, 77: 7–23.

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

Galor, Oded. 2005. "From Stagnation to Growth: Unified Growth Theory." In *Handbook of Economic Growth.* , ed. P. Aghion and S. N. Durlauf, 171–293. North-Holland.

Galor, Oded. 2011. *Unified Growth Theory.* Princeton: Princeton University Press.

Galor, Oded, and David N. Weil. 2000. "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond." *American Economic Review*, 90: 806–828.

Galor, Oded, and Omer Moav. 2006. "Das Human-Kapital: A Theory of the Demise of the Class Structure." *Review of Economic Studies*, 73: 85–117.

Georgescu-Roegen, Nicholas. 1971. *The Entropy Law and the Economic Process.* Harvard University Press.

Georgescu-Roegen, Nicholas. 1975. "Energy and Economic Myths." *Southern Economic Journal*, 41: 347–381.

Gillings, Michael R., Martin Hilbert, and Darrell J. Kemp. 2016. "Information in the Biosphere: Biological and Digital Worlds." *Trends in Ecology and Evolution*, 31: 180–189.

Gomułka, Stanisław. 2023. *The Global Long-term Economic Growth and the Economic Transformation of Poland and Eastern Europe.* Scholar Publishing House.

Gordon, Robert J. 1990. *The Measurement of Durable Goods Prices.* University of Chicago Press.

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

Grace, Katja. 2013. "Algorithmic Progress in Six Domains." Berkeley, CA: Machine Intelligence Research Institute Technical report 2013-3.

Grace, Katja, Harlan Stewart, Julia F. Sandkuehler, Stephen Thomas, Ben Weinstein-Raun, and Jan Brauner. 2024. "Thousands of AI Authors on the Future of AI." arXiv 2401.02843.

Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." *Review of Economics and Statistics*, 100: 753–768.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. 1997. "Long-Run Implications of Investment-Specific Technological Change." *American Economic Review*, 87: 342–362.

Growiec, Jakub. 2007. "Beyond the Linearity Critique: The Knife-Edge Assumption of Steady-State Growth." *Economic Theory*, 31: 489–499.

Growiec, Jakub. 2013. "A Microfoundation for Normalized CES Production Functions with Factor-Augmenting Technical Change." *Journal of Economic Dynamics and Control*, 37: 2336–2350.

Growiec, Jakub. 2018. "Factor-Specific Technology Choice." *Journal of Mathematical Economics*, 77: 1–14.

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

Growiec, Jakub. 2022*b*. "Automation, Partial and Full." *Macroeconomic Dynamics*, 26: 1731–1755.

Growiec, Jakub. 2022*c*. “R&D Capital: An Engine of Growth.” *Economics Letters*, 217: 110703.

Growiec, Jakub. 2023*a*. “Industry 4.0? Framing the Digital Revolution and Its Long-Run Growth Consequences.” *Gospodarka Narodowa. The Polish Journal of Economics*, 4: 1–16.

Growiec, Jakub. 2023*b*. “What Will Drive Global Economic Growth in the Digital Age?” *Studies in Nonlinear Dynamics and Econometrics*, 27: 335–354.

Growiec, Jakub, and Jakub Mućk. 2020. “Isoelastic Elasticity of Substitution Production Functions.” *Macroeconomic Dynamics*, 24(7): 1597–1634.

Growiec, Jakub, Peter McAdam, and Jakub Mućk. 2023. “R&D Capital and the Idea Production Function.” Federal Reserve Bank of Kansas City Working paper RWP 23-05.

Ha, Joonkyung, and Peter Howitt. 2007. “Accounting for Trends in Productivity and R&D: A Schumpeterian Critique of Semi-Endogenous Growth Theory.” *Journal of Money, Credit and Banking*, 39(4): 733–774.

Hanson, Robin. 2000. “Long-Term Growth as a Sequence of Exponential Modes.” George Mason University Working paper.

Hanson, Robin. 2001. “Economic Growth Given Machine Intelligence.” University of California, Berkeley Working paper.

Hanson, Robin, and Eliezer Yudkowsky. 2013. *The Hanson-Yudkowsky AI-Foom Debate.* Machine Intelligence Research Institute.

Harari, Yuval Noah. 2017. *Homo Deus: A Brief History of Tomorrow.* Vintage.

Hemous, David, and Morten Olsen. 2018. “The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality.” University of Zurich Working paper.

Henderson, David J., and R. Robert Russell. 2005. “Human Capital and Convergence: A Production–Frontier Approach.” *International Economic Review*, 46: 1167–1205.

Hercowitz, Zvi. 1998. “The ‘Embodiment’ Controversy: A Review Essay.” *Journal of Monetary Economics*, 41: 217–224.

Hernandez, Danny, and Tom B. Brown. 2020. “Measuring the Algorithmic Efficiency of Neural Networks.” arXiv Preprint arXiv:2005.04305.

Hilbert, Martin, and Priscila López. 2011. "The World's Technological Capacity to Store, Communicate, and Compute Information." *Science*, 332: 60–65.

Ho, Anson, Tamay Besiroglu, Ege Erdil, David Owen, Robi Rahman, Zifan Carl Guo, David Atkinson, Neil Thompson, and Jaime Sevilla. 2024. "Algorithmic Progress in Language Models." arXiv 2403.05812.

Jones, Charles I. 1995. "R&D-Based Models of Economic Growth." *Journal of Political Economy*, 103: 759–84.

Jones, Charles I. 1999. "Growth: With or Without Scale Effects?" *American Economic Review*, 89(2): 139–144.

Jones, Charles I. 2002. "Sources of U.S. Economic Growth in a World of Ideas." *American Economic Review*, 92: 220–239.

Jones, Charles I. 2005*a*. "Growth and Ideas." In *Handbook of Economic Growth.* , ed. P. Aghion and S.N. Durlauf. North-Holland.

Jones, Charles I. 2005*b*. "The Shape of Production Functions and the Direction of Technical Change." *Quarterly Journal of Economics*, 120: 517–549.

Jones, Charles I., and Christopher Tonetti. 2020. "Nonrivalry and the Economics of Data." *American Economic Review*, 110: 2819–2858.

Jones, Charles I., and Jihee Kim. 2018. "A Schumpeterian Model of Top Income Inequality." *Journal of Political Economy*, 126: 1785–1826.

Jones, Charles I., and Paul M. Romer. 2010. "The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital." *American Economic Journal: Macroeconomics*, 2: 224–245.

Jones, Larry E., and Rodolfo E. Manuelli. 1990. "A Convex Model of Equilibrium Growth: Theory and Policy Implications." *Journal of Political Economy*, 98: 1008–1038.

Jorgenson, Dale W. 1995. *Productivity. Volume 1: Postwar U.S. Economic Growth.* MIT Press.

Jorgenson, Dale W. 2005. "Accounting for Growth in the Information Age." In *Handbook of Economic Growth.* , ed. P. Aghion and S.N. Durlauf. North-Holland.

Jumper, John, Richard Evans, Alexander Pritzel, et al. 2021. "Highly Accurate Protein Structure Prediction with AlphaFold." *Nature*, 596: 583–589.

Kaldor, Nicholas. 1961. "Capital Accumulation and Economic Growth." In *The Theory of Capital.* , ed. L. A. Lutz and D. C. Hague, 177–222. Palgrave Macmillan.

Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *Quarterly Journal of Economics*, 129(1): 61–103.

Klein, Lawrence R., and Richard F. Kosobud. 1961. "Some Econometrics of Growth: Great Ratios of Economics." *Quarterly Journal of Economics*, 75: 173–198.

Klump, Rainer, Peter McAdam, and Alpo Willman. 2007. "Factor Substitution and Factor Augmenting Technical Progress in the US." *Review of Economics and Statistics*, 89: 183–192.

Klump, Rainer, Peter McAdam, and Alpo Willman. 2012. "Normalization in CES Production Functions: Theory and Empirics." *Journal of Economic Surveys*, 26: 769–799.

Koop, Gary, Jacek Osiewalski, and Mark F. J. Steel. 1999. "The Components of Output Growth: A Stochastic Frontier Analysis." *Oxford Bulletin of Economics and Statistics*, 61: 455–487.

Koop, Gary, Jacek Osiewalski, and Mark F. J. Steel. 2000. "Measuring the Sources of Output Growth in a Panel of Countries." *Journal of Business and Economic Statistics*, 18: 284–299.

Korinek, Anton, and Donghyun Suh. 2024. "Scenarios for the Transition to AGI." University of Virginia and Brookings.

Korinek, Anton, and Joseph Stiglitz. 2019. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." In *The Economics of Artificial Intelligence: An Agenda.* , ed. A. Agrawal, J. S. Gans and A. Goldfarb. University of Chicago Press.

Korinek, Anton, and Megan Juelfs. 2022. "Preparing for the (Non-Existent?) Future of Work." In *The Oxford Handbook of AI Governance*. Oxford University Press.

Kremer, Michael. 1993. "Population Growth and Technological Change: One Million B.C. to 1990." *Quarterly Journal of Economics*, 108: 681–716.

Kruse-Andersen, Peter K. 2023. "Testing R&D-Based Endogenous Growth Models." *Oxford Bulletin of Economics and Statistics*, forthcoming (Online First).

Krusell, Per, Lee E. Ohanian, José-Victor Ríos-Rull, and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 68: 1029–1054.

Kumar, Subodh, and R. Robert Russell. 2002. "Technological Change, Technological Catch-up, and Capital Deepening: Relative Contributions to Growth and Convergence." *American Economic Review*, 92: 527–548.

Kurzweil, Ray. 2005. *The Singularity is Near.* New York: Penguin.

Lewandowski, Piotr, Albert Park, Wojciech Hardy, Yang Du, and Saier Wu. 2022. "Technology, Skills, and Globalization: Explaining International Differences in Routine and Nonroutine Work Using Survey Data." *World Bank Economic Review*, 36: 687–708.

Madsen, Jakob. 2008. "Semi-endogenous Versus Schumpeterian Growth Models: Testing the Knowledge Production Function Using International Data." *Journal of Economic Growth*, 13: 1–26.

McAdam, Peter, and Alpo Willman. 2018. "Unraveling the Skill Premium." *Macroeconomic Dynamics*, 22: 33–62.

Milanović, Branko. 2016. *Global Inequality: A New Approach for the Age of Globalization.* Harvard University Press.

Mućk, Jakub. 2017. "Elasticity of Substitution Between Labor and Capital: Robust Evidence from Developed Economies." Narodowy Bank Polski Working paper No. 271.

Nordhaus, William D. 2021. "Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth." *American Economic Journal: Macroeconomics*, 13: 299–332.

OpenAI. 2023. "GPT-4 Technical Report." arXiv 2303.08774.

Piketty, Thomas. 2014. *Capital in the Twenty-First Century.* Harvard University Press. Translated from French by Arthur Goldhammer.

Piketty, Thomas, and Gabriel Zucman. 2014. "Capital is Back: Wealth-Income Ratios in Rich Countries, 1700-2010." *Quarterly Journal of Economics*, 129(3): 1255–1310.

Romer, Paul M. 1986. "Increasing Returns and Long-run Growth." *Journal of Political Economy*, 94(5): 1002–1037.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98: S71–S102.

Roodman, David. 2020. "On the Probability Distribution of Long-Term Changes in the Growth Rate of the Global Economy: An Outside View." Open Philanthropy Report.

Roser, Max, Hannah Ritchie, and Edouard Mathieu. 2023. "What is Moore's Law?" *Our World in Data*. https://ourworldindata.org/moores-law.

Sachs, Jeffrey D., Seth G. Benzell, and Guillermo LaGarda. 2015. "Robots: Curse or Blessing? A Basic Framework." NBER Working Paper Np. 21091.

Sevilla, Jaime, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos. 2022. "Compute Trends Across Three Eras of Machine Learning." arXiv 2202.05924.

Silver, David, Thomas Hubert, Julian Schrittwieser, et al. 2018. "A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go Through Self-Play." *Science*, 362: 1140–1144.

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

Solow, Robert M. 1957. "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics*, 39: 312–320.

Solow, Robert M. 1960. "Investment and Technical Progress." In *Mathematical Methods in the Social Sciences.* , ed. K. Arrow, S. Karlin and P. Suppes, 89–104. Stanford University Press.

Spitz-Oener, Alexandra. 2006. "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure." *Journal of Labor Economics*, 24: 235–270.

Tegmark, Max. 2017. *Life 3.0: Being Human in the Age of Artificial Intelligence.* New York: Knopf.

Temple, Jonathan. 2006. "Aggregate Production Functions and Growth Economics." *International Review of Applied Economics*, 20: 301–317.

Timmer, Marcel P., and Bart van Ark. 2005. "IT in the European Union: A Driver of Productivity Divergence?" *Oxford Economic Papers*, 57(4): 693–716.

Trammell, Philip, and Anton Korinek. 2021. "Economic Growth under Transformative AI: A Guide to the Vast Range of Possibilities for Output Growth, Wages, and the Labor Share." Oxford University Global Priorities Institute WP 8-2020.

Uzawa, Hirofumi. 1961. "Neutral Inventions and the Stability of Growth Equilibrium." *Review of Economic Studies*, 28: 117–124.

Waldrop, M. Mitchell. 2016. "The Chips Are Down for Moore's Law." *Nature*, 530: 144–147.

Yudkowsky, Eliezer. 2013. "Intelligence Explosion Microeconomics." Machine Intelligence Research Institute Technical report 2013-1.

Zeira, Joseph. 1998. "Workers, Machines, and Economic Growth." *Quarterly Journal of Economics*, 113: 1091–1117.

The Center on Regulation and Markets at Brookings provides independent, non-partisan research on regulatory policy, applied broadly across microeconomic fields. It creates and promotes independent economic scholarship to inform regulatory policymaking, the regulatory process, and the efficient and equitable functioning of economic markets.

Questions about the research? Email communications@brookings.edu.
Be sure to include the title of this paper in your inquiry.