Extended outline for submission to NBER Conference "AI and Economic Measurement", Stanford, CA — May 7, 2026

Is Software Eating the World?

Measuring the Progress and Diffusion of AI*

Filippo Bontadini, LUISS University and SPRU, University of Sussex
Carol Corrado, Center for Business and Public Policy, McDonough School of Business, Georgetown University
Jonathan Haskel, Imperial Business School, Imperial College London
Cecilia Jona-Lasinio, LUISS Business School

This draft: January 2026

Abstract

How can we trace AI in the national accounts? What changes need to be made to track it better? This outline note sets out a framework for measuring some of the key elements of AI technology. It comments on the existing evidence, in particular on factor rental shares and productivity, and suggests future developments that would improve such evidence. We show initial evidence for the US and cross-countries.

Keywords: Artificial intelligence (AI), innovation, productivity growth, intangibles
JEL reference: O47, E22, E01

*Contacts: Carol Corrado carol.corrado@georgetown.edu. Jonathan Haskel j.haskel@imperial.ac.uk. Cecilia Jona-Lasinio <cjonalasinio@luiss.it>; Filippo Bontadini fbontadini@luiss.it. None of the authors have financial or other ties to organisations related to this research.

1. Introduction

This paper will aim to address some of the emerging measurement challenges of quantifying AI's effects on the economy. The central goal is to convert AI into measurable economic terms, either aligning with current national accounts or suggesting new approaches.

One reason why this is important is that theoretical predictions of the effect of AI are, thus far, very wide-ranging. Regarding productivity, The Federal Reserve Bank of Dallas has a now-famous chart showing either unlimited growth or utter collapse (https://www.dallasfed.org/research/economics/2025/0624). Likewise, predictions of the evolution of the share of capital returns vary from no change to 100%.

That said, theory does give some indication of the evidence we shall need to understand the effects of AI. For example, much modelling crucially depends on whether the elasticity of substitution between labour and (AI-powered) capital is more or less than unity, indications of which would appear in factor shares of the relevant inputs to AI.[1] Similarly, the widely-discussed phenomenon of software that writes its own software are, like the old literature on machines producing machines, amenable to measurement if we can identify the use of software in the production of software.

What existing evidence, if any, do we have to bear on these predictions? One strand of literature is to look at tasks to try to understand what tasks AI might replace. This is limited in that we need a number of assumptions to graft tasks to national accounts indicators, and AI might replace existing tasks with new ones. Another strand is to look at the experience of specific worker types e.g. software programmers. This is limited to that specific group. Further, neither of these approaches speak directly to what might happen to the rate of return or capital share.

This research will try a different route. We start by asking: how might we measure AI and identify its effects in the national accounts, at least as currently constituted? If we can do that, we then move to reviewing the evidence we have so far on AI's effects.

On the first question, as has been pointed out, a underlying technology of the current wave of AI is machine learning. If we think of this as (a) powerful computers ("compute") (b) running improving software (c) scanning large databases, then we have a ready relation to the national accounts. Computers are a tangible asset, and software and databases intangible assets. Computers and software are currently measured as assets, with databases partially integrated in

[1] The factor share of input X is of particular interest since this has been used as a diagnostic for the elasticity of substitution between input X and other factors. Under the "bottlenecks" style assumption, AI-powered tasks have an elasticity of substitution of less then one with labour. In this case, as AI undertakes more tasks and becomes cheaper it's factor payments share falls owing to the limited rise in quantities due to the relatively low elasticity of substitution.

software, but to be fully implemented in the future. Conceptualised in this way suggests looking at investment in hardware and software. Further, if we can identify the software-producing industry, we might look at the use of hardware and software to produce new software. If we can measure the capital stock and rental prices, we may look at the shares of value added accounted for by factor payments to hardware and software, to shed light on possible degrees of substitution. Finally, to measure AI's effect on productivity we may look at the contributions to labour productivity of software production and software use (the latter being software capital adoption across industries).

Measurement of all these elements is however very complicated. Consider for example the factor payments share of software, which needs the stock of software, its (rental) price and the value added of which it is a share. In turn, the rental rate requires knowledge of the rate of return on capital. The rate of return on capital itself depends on what capital is measured: as Corrado et al, 2022 point out, with intangible assets being of growing importance, omitting intangibles biases measurement of the rate of return.

This paper then presents some results but sets out what might be measured in the future. It starts by justifying looking at software and hardware as an indicator of AI. We then measure the evolution of its factor share and the attendant rates of return. Finally, we look at early indicators of the contribution of software to productivity growth.

2. AI, software, hardware, databases and "compute"

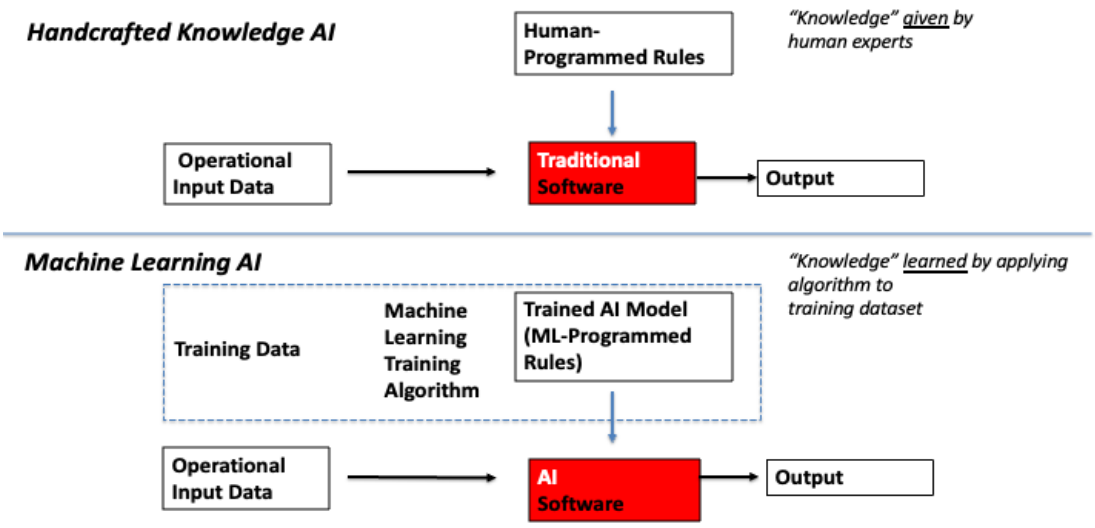
In seeking to explain what is different about the modern wave of AI, Allen (2020) distinguishes between two AI types, "Handcrafted Knowledge" AI and "Machine Learning" AI. Figure 1 upper panel shows "handcrafted" AI. The AI chess-playing system Deep Blue, which defeated the reigning world chess champion Garry Kasparov in 1997, was developed through a collaboration between computer programmers and human chess grandmasters. The "knowledge" of the computer was the hand-typed computer code that told the computer the rules but also the potential moves and countermoves written in the light of what grandmasters thought was winning play.

The lower panel shows Machine Learning AI. Here, knowledge is not programmed by humans. Rather, the knowledge is learned by the computer from data. The initial impetus is a training dataset on which a ML algorithm runs, producing then an AI model. This model is the basis of AI software into which the user inputs their operational data to be processed for output.

The more complex and computationally intensive generative AI applications where the computer generates answers to queries are called "deep learning," which may be thought of as a subset of machine learning. Facial recognition is another example of the complexity computers can now handle. The recognition of images, which humans can do easily, was

previously too hard to program in Handcrafted style. But with ML/DL the computer, once trained, can be used to recognize faces and images.

Figure 1: Software creation pre- and post-machine learning



Source. Authors' adaptation of Allen (2020).

Baily et al. (2025) provide a history of advancements in AI. A conference held in 1956 at Dartmouth spurred early research on AI, which concentrated on classification applications rather than generative AI or question answering. Progress was constrained by the limited computational power available at the time. By the 2010s systems such as Alexa and Siri had been introduced, and the large technology firms were using machine learning algorithms to forecast demand. Computational power and near-free storage made this era popularly described as one of "Big Data."

A significant turning point in artificial intelligence research occurred in 2012, initiating the deep learning revolution. The AlexNet model, developed by University of Toronto researchers, achieved a decisive victory at the ImageNet Large Scale Visual Recognition Challenge in that year, thereby demonstrating the effectiveness of deep learning techniques for image recognition tasks. By this time large language models (LLMs) had emerged, and newspapers were generating some stories by machine, but only for well-defined tasks, like reporting sports scores, or financial data. The breakthrough in neural architecture in 2017 (the publication of the transformer architecture paper by the Google Brain researchers, Vaswani et al, 2017) greatly improved the speed and accuracy of language processing in LLMs, kicking off the current era of generative AI.

In sum, the lower panel of Figure 1 suggests it seems reasonable to think of the production of AI tools from a combination of software, hardware/compute and databases. The use of such tools would be the use of software in the wider economy. Measuring "data centres" might also suggest including buildings in the computer services industry (in national accounts terms these would be in "other structures"). "Compute" is generally taken to refer to hardware use in the generation of AI models. "Cloud services" is the purchase of computer services, often comprising both hardware and software with minimal hardware and software usage at the purchasing firm from the computer services industry. Thus we might capture this as growing provision of computer hardware and software capital services by the computer services industry. "Software writing software" would be software capital deepening in the software writing industry. Finally, R&D is problematic, if the conduct of R&D requires writing software. As we review, the European convention is to include this in the software data, but in the US in the R&D data.

3. Measurement framework

We set out here the measurement framework for capital assets to try to be clear what we measure below (e.g. capital payment shares) and what is needed to construct these data.

We wish to compare as far as possible across countries. In the below, we denote country, *c*, industry, *j*, at time *t*. Capital assets, which will be both tangible and intangible are denoted *a*. Labour types are *b*. In some of the below, we ignore country *c* for clarity

1.1 Basic approach: capital shares

1.1.1 Capital assets

We start with nominal investment data in asset *a* for industry j at time t, which we deflate by an investment deflator; we write this as

$$I_{a,j,t} = \frac{P_{I,a,j,t} I_{a,j,t}}{P_{I,a,j,t}} \qquad (1)$$

As we discuss below, our investment price indices are harmonised across countries for some assets. Equation (1) is written in this way to make the point that the investment deflator might not be the "true" deflator if it is not quality-adjusted etc.

We use this real investment to build a real capital stock for industry *j*, capital asset *a* at time *t* via the perpetual inventory method whereby for any capital asset the stock of that asset evolves according to

$$K_{a,j,t} = I_{a,j,t} + (1-\delta_{a,j,t})K_{a,j,t-1} \quad (2)$$

Where, I is real investment over the relevant period and δ the geometric rate of depreciation. As a practical matter, the starting value for this equation varies by asset, with tangible asset starting values being taken from previous versions of the EUKLEMS dataset.

The investment price is converted into a rental price using the Hall-Jorgenson relation ((Hall and Jorgenson 1967), where we assume an industry-specific net rate of return such that the capital rental price times the capital stock equals the total industry operating surplus (on all of this, see for example, Oulton and Srinivasan (2003)).

Thus rental prices are given by

$$P_{Ka.,j,t} = \tau_{Ka,j,t}(r_{j,t} + \delta_{a,j,t} - \pi_{Ka,j,t})P_{Ia,j,t} \quad (3)$$

where $\tau_{Ka,j,t}$ is an asset-industry-time specific tax-adjustment factor, $\pi_{Ka,j,t}$ is a capital gains term and $P_{Ia,j,t}$ is the investment price of asset *a*.

1.1.2 Capital input rental cost share

With a capital stock K, we cannot calculate Pk without knowing r. A number of assumptions are possible, and we shall explore the consequences of relaxing them. Our central estimate follows Jorgenson/Griliches and assumes a common rate of return across assets within industries. Further, we have nominal value added measures by industry. Thus, we can measure returns to capital residually if we assume

$$\sum_a P_{Ka,j,t}K_{a,j,t} = P_{Q,j,t}Q_{j,t} - P_{L,j,t}L_{j,t} \quad (4)$$

Where the final term is the wage bill. We compute r so that this equality holds. We define what we shall call an "input rental share" (as opposed to an output value-added share), for capital asset *a*

$$w_{Qj,t}^{Ka,j,t} = \frac{P_{Ka,j,t}K_{a,j,t}}{P_{Q,j,t}Q_{j,t}} \quad (5)$$

1.2 Productivity

1.2.1 Industry outputs and inputs

Industry j uses capital and labour services (respectively K_j and L_j) to produce (value-added) output Q_j. We denote value added by Q to stress that value added includes intangible capital investment. Thus, for each industry, we have the implicitly value-added defined Total Factor Productivity (TFP), $\Delta \ln TFP_j$:

$$\Delta \ln Q_{j,t} \equiv \overline{w_{Qj.t}^{Ka,j,t}} \Delta \ln K_{a,j,t} + \overline{w_{Qj,t}^{Lb.j,t}} \Delta \ln L_{b,j,t} + \Delta \ln TFP_{j,t}$$

$$\textit{where} \qquad (6)$$

$$w_{Qj,t}^{Ka,j,t} = \frac{\sum_a P_{Ka,j,t} K_{a,j,t}}{P_{Q,j} Q_j}, w_{Qj}^{Lb,j,t} = \frac{\sum_b P_{Lb,j,t} H_{b,j,t}}{P_{Q,j} Q_j}, \overline{w} = 0.5 * (w_t + w_{t-1})$$

Where the terms in "w" are shares of factor rental costs in industry nominal value-added, averaged over two periods and K and L are for each asset (or in the case of labour, each worker type b). TFP growth is calculated as a residual.

1.3 Aggregation across industries and sources of economy-wide growth

Define changes in *aggregate* real value added as a weighted sum of changes in industry real value added, where the weights are nominal industry value-added as a share of aggregate value-added: we shall denote these with a "v" since they are industry value-added as a share of total value added (the "w" weights are rental payment shares)

$$\Delta \ln Q \equiv \sum_j \overline{v_Q^{Qj}} \Delta \ln Q_j$$

$$\textit{where} \qquad (7)$$

$$v_Q^{Qj} = \frac{P_{Qj} Q_j}{\sum_j P_{Qj} Q_j}, \bar{v} = 0.5 * (v_t + v_{t-1})$$

We may aggregate hours similarly (notice this drives a wedge between labour productivity defined in this "bottom up" way and labour productivity defined as aggregate ΔlnV, a value-added weighted sum of industry $\Delta \ln V_j$ and naturally summed hours). Ignoring this and assuming constant returns, the relation between aggregate labour productivity growth (where

hours are aggregated as weighted logarithmic changes) and industry-level inputs can be written as:

$$\Delta \ln(Q/H)_{c,t} = \sum_{c,j} \overline{\left(\frac{\sum_a P_{Kc,a,j,t} K_{c,a,j,t}}{P_{Qc,j,t} Q_{c,j,t}} \right)} \Delta \ln(K/H)_{c,a,j,t} + \sum_{c,j} \overline{\left(\frac{\sum_b P_{Lc,b,j,t} H_{c,b,j,t}}{P_{Qc,j,t} Q_{c,j,t}} \right)} \Delta \ln(L/H)_{c,b,j,t} + \sum_{c,j} \overline{\left(\frac{P_{Qj} Q_j}{\sum_j P_{Qj} Q_j} \right)} \Delta \ln TFP_{c,j,t} \quad (6)$$

Where the weights are expressed as value-added shares, the bar over the weight terms indicates a two period average and we introduce the country dimension *c*.

This means the contribution of asset *a* to labour productivity in the whole economy is the extent of its capital deepening weighted by its rental payment, plus the TFP gain in the industry producing that asset weighted by its value added. We write this for the whole economy.

$$\Delta \ln(Q/H)^a_{c,t} = \overline{\left(\frac{P_{Kc,a,j,t} K_{c,a,j,t}}{P_{Qc,j,t} Q_{c,j,t}} \right)} \Delta \ln(K/H)_{c,a,j,t} + \overline{\left(\frac{P_{Qj} Q_j}{\sum_j P_{Qj} Q_j} \right)^a} \Delta \ln TFP^a{}_{c,j,t} \quad (8)$$

We shall also look at capital deepening within the sector producing software. The idea of "software writing itself" could be captured as software capital deepening i.e. software substituting for labour in the software-producing industry. The impact of data centres would be seen, in this framework, via structures, hardware and software capital deepening within the software service providing industry.

1.4 Relation with the task literature: Interpreting technical change and capital deepening

This framework is broadly applicable and measurable, but it raises the question of the determinants of the technical change terms. Following endogenous growth theories, total factor productivity as the Hicksian production function "shifter" is a "costless" source of growth that results from the nonrival character of intangible capital, i.e., TFP represents knowledge spillovers, or diffusion. For new AI software (or AI-enabled machines), one "costless" source-of-growth is open-source software code, which according to Calderon et al (2022) grew more than 10 percent per year from 2010 to 2019; (see also Hoffman, Nagle and Zhou 2024).

Thus, while some factors affecting the diffusion of technology in this model can be identified, the model is not very specific about the underlying mechanisms that drive technological diffusion via the ΔlnTFP. By contrast, the task approach frames technical change in terms of specific mechanisms. Tasks are set as the fundamental unit of production in this approach, where tasks are performed by unskilled labor, skilled labor and capital (e.g., machines

and software) in various combinations. By parameterizing these combinations via certain elasticities, e.g., capital-labor substitution, the task model is amenable to considering displacement of tasks by automation as well as reinstatement via new tasks (i.e., augmentation), where the impact of augmentation varies according to complementarities and market forces affecting labor supply. This modelling is a sharp departure from traditional skill-biased technical change models, which usually assume augmentation works as a uniform boost to productivity of all tasks that labor performs.

We believe our growth accounting approach complements the rich bottom-up nature of the task framework even though the underlying task reallocation and its separate processes (automation and augmentation) are not visible. As described above, automation is a driver of capital deepening, while the pace of augmentation moderates the ultimate impact of AI on job loss. The task model also suggests that augmentation will be greater or smaller depending on the extent to which technology complements labour, i.e., the extent to which labour has been made more productive overall. When taking (7) to the data, and thinking in terms of the task model, measured ΔlnTFP then reflects the net savings from both the displacement and reinstatement of workers due to technical change. In the classic interpretation of the microeconomics of change due to a GPT, as Bresnahan (2021) explains, TFP rather reflects the net impact of the cost savings due to the organizational change that complements the technology. Our approach is congruent with both these views of the underlying forces at work in the economy.

The model assumes constant returns and competitive factor payments. Using (7) to estimate the impact of a new technology on productivity rests on a foundational result in productivity and growth accounting due to Hulten that predicts, without knowing the full structure of the economy, the first order macro effects of micro/sectoral shocks under constant returns and competition (Hulten 1978). Theoretically, nonlinearities (e.g., complementarities), network effects (which tend to manifest as increasing returns), and misallocation can alter the aggregate impacts implied by Hulten's theorem (Baqaee and Farhi 2019). In their work using the task approach to estimate the aggregate impact of AI, Filippucci et al (2024) examined the impact of these distortionary influences on the applicability of Hulten's theorem. After including the terms Baqaee and Farhi (2019) suggest are needed when non-linearities and barriers to reallocation are present, in most cases, the impact of additional terms turned out to be relatively small compared with the first-order effect; (see their Table 2).

Finally, for simplicity we have not set out here different sectors. Oulton sets out, discussed more below, a two sector model with ICT production and consumption goods production, where the latter uses ICT. In that model, the diffusion of ICT in the economy depends on the adoption of ICT capital in the consumption sector, the pace of which depends on the relative pace of TFP growth in both sectors. If ICT TFP growth is faster than

consumption sector TFP growth, then the relative price of ICT goods falls relative to consumption goods and ICT capital is increasingly adopted. Thus, labour productivity growth depends on both TFP in the consumption sector and the additional TFP in the ICT sector. This suggests looking at difference in TFP growth and prices between AI-producing and AI-using sectors, which we hope to take up in future work.

4. US data on income shares

We start with the US. For the US we have official national accounts data 1987-2024. These data are then more timely than the cross-country data we look at below, which are 1995- 2021. They capitalise a number of intangible assets (software, R&D, artistic originals): in the cross-country data we shall capitalise a broader range.

One important adjustment we shall make to these data concerns R&D. In recent years much R&D is conducted in the software-producing sector and much R&D conducted anywhere consists of writing software (see Bontadini et al, 2026 for a discussion). Thus we look at three industries and three rental share indicators. The industries are

i. the "whole economy" (non-farm business);
ii. the "software producing" sector; Publishing industries, except internet (includes software); Data processing, internet publishing, and other information services; Computer systems design and related services.
iii. The "whole economy, excluding software producing" i.e. (i) excluding (ii).

Our use here of the software-producing industries, essentially the main part of what is called "ICT services", is a first pass at a difficult measurement problem. US data suggests around 20% of software is written on the own-account and therefore produced in the host industry e.g. in financial services if own-account banking software. Thus our industries understate software creation. Against this, our industries do more than software provision e.g. the publishing industries include newspapers. Further investigation here is warranted.

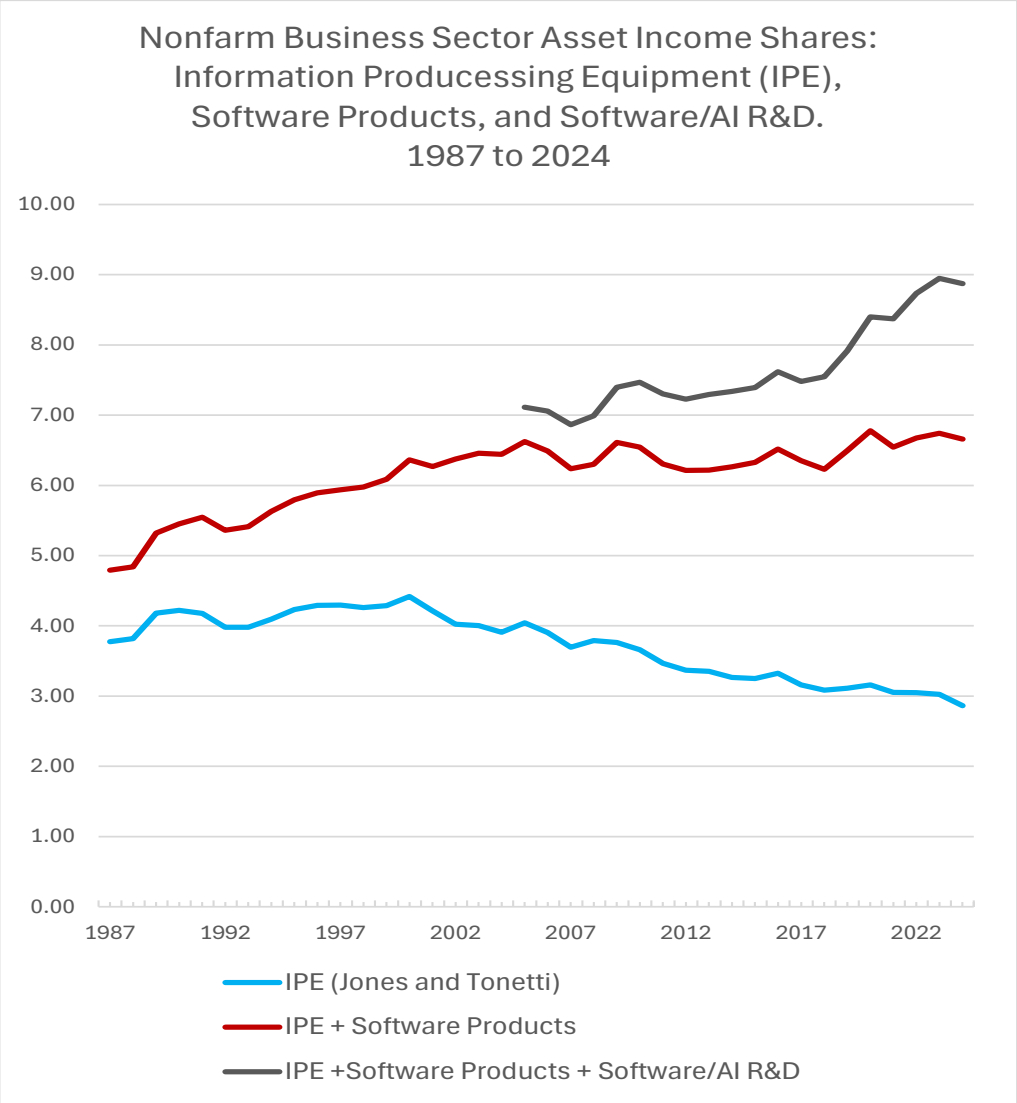
Our asset rental share indicators are factor payments for the following assets

i. IPE (Information processing equipment/hardware), consisting of Computers and peripheral equipment, communication equipment, medical equipment and instruments, non-medical instruments, photo copy and related equipment, office an accounting equipment.
ii. Software products; consisting of rental payments to software used as an input in the economy (except software produced in the course of conducting R&D)

iii. Software/AI R&D; consisting of rental payments to software produced in the course of conducting R&D).

Figure shows our results for share in value added of the US non-farm business sector of asset rental payments as above.

Figure 2: Rental income shares, non-farm business economy



Source: author calculations from BLS multi-factor productivity data

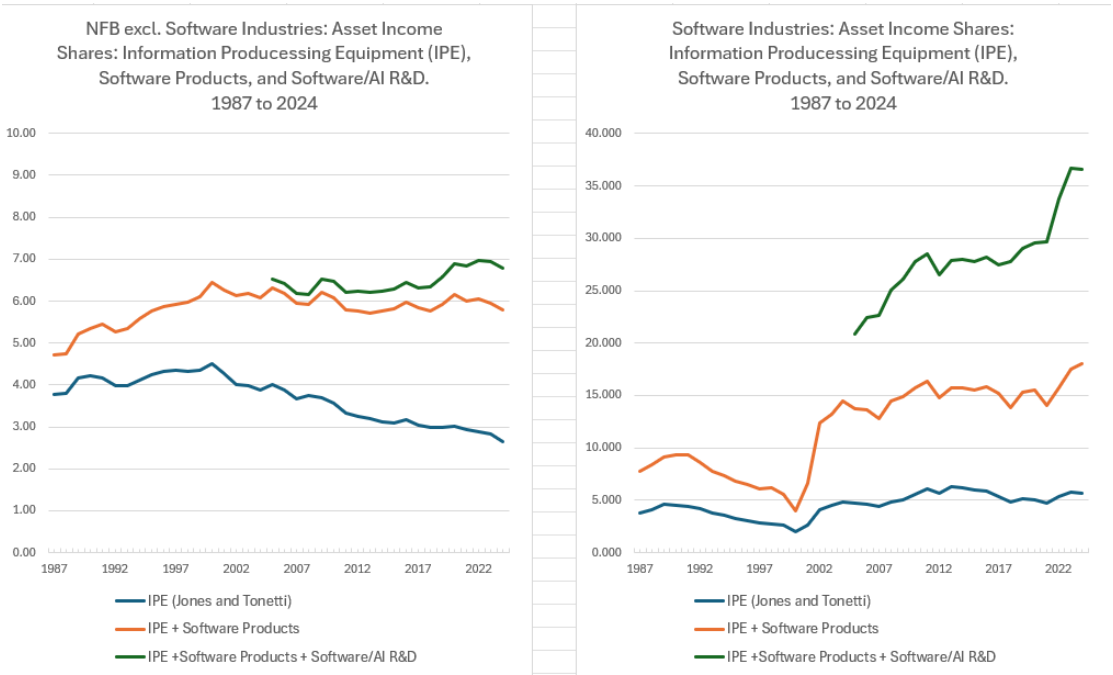
As the graph shows, the IPE/hardware share is falling (the lowest line). However, if we add in payments to software, the rental shares are rising. This is especially true if we add in the software produced in the course of conducting R&D (the uppermost line).

Figure shows the same data but for the non-farm business sector excluding software producing industries and for the software-producing industries. The differences are of interest. The left panel shows a similar picture to Figure , namely a falling IPE/hardware share and rising, then flat software shares. It is interesting to note that excluding the software-producing sector is sufficient to flatten out the software share. It is also of interest to note the post 2017 rise in the share including R&D software, a rise just after the transformer architecture took hold. This suggests that this approach seems to be picking up some elements of the AI revolution.

The right panel shows the software-producing industries, where here the software/AI R&D is the payments to software created in the course of undertaking R&D in this industry. We have some interesting findings.

First, notice that all the shares rise in the early 2000s. This could reflect the transition to software as a service (Salesforce launched its services in 1999 for example). Second, in recent years there is a strong rise in payments to software including those produced in the course of R&D, in line with the advances in AI.

Figure 3: Rental income shares, non-farm business economy



Source: author calculations from BLS multi-factor productivity data

5. Cross country data on income shares

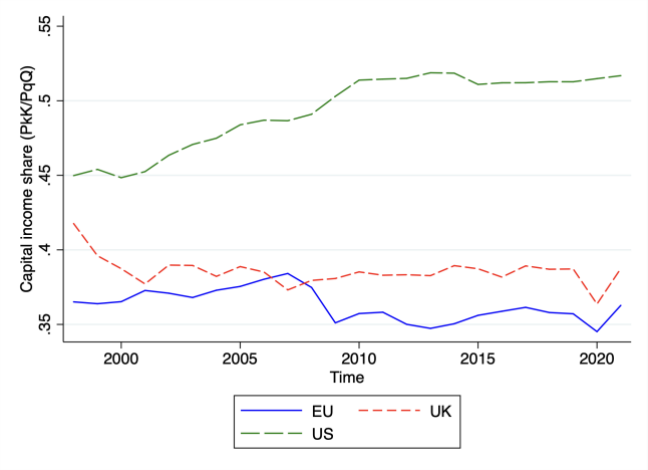
To make cross-country comparisons, we use the EUKLEMS & INTANProd productivity dataset (LLEE 2025). We draw upon data from its analytical module (described in Bontadini et al. 2023) for the European Union, measured as an aggregate of 10 member countries for which we have both full growth accounts and substantial history (Austria, Czechia, Denmark, Finland, France, Germany, Italy, Netherlands, Spain, and Sweden), plus the United Kingdom and the United States. The EUKLEMS & INTANProd data will be updated shortly, allowing for this cross-country comparison to be extended up to 2023. For the moment, our data end in 2021. Concerning industries, in the cross-country work we exclude real estate, agriculture, health, education and defence. The software producing sector is not as closely defined as in the US-only work above. For the US it includes: NAICS Publishing industries, except internet (includes software) 5110; Data processing, internet publishing, and other information services 5140; Computer systems design and related services 5415. For the EU countries it refers to NACE Rev 2 sectors J58-J60 Publishing, audio-visual and broadcasting activities and J62-J63 IT and other information services.

4.1 Factor shares and rates of return

Figures 4 and 7 show our results for factor shares. Starting from the broad non-farm business sector, Figure 2 shows a significant divergence in capital income shares between the United States, UK and the EU. The US has experienced a sustained increase in the aggregate capital income share, rising from approximately 0.45 in the late 1990s to over 0.50 by 2021. ,. Conversely, in the European countries and the UK capital shares remain consistently below the US levels, possibly reflecting a different pace of AI technological adoption.

Figure 3: –US, UK and EU, Total assets - Capital income shares

Non-farm business sector

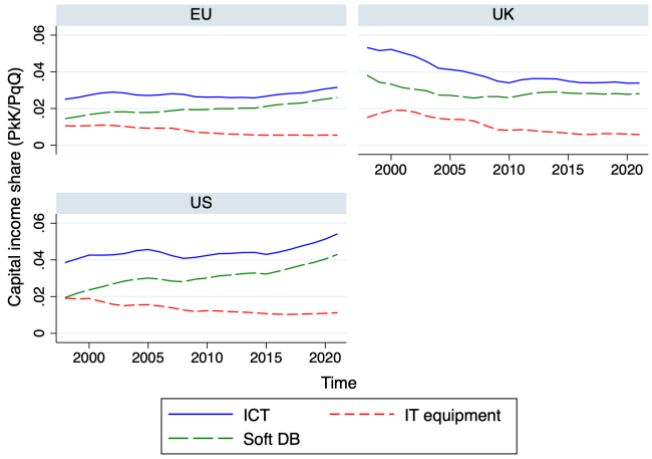


Source: author calculations from EUKLEMS and INTANProd database.

These aggregate dynamics are further explored in Figure 5, which decomposes ICT capital income shares into “Hardware” (IT equipment), and "Software" (software and databases). These shares are shown individually, not cumulated as in the US figures above. Across the US, UK and the EU there is a clear decline in the income share of tangible ICT hardware. Simultaneously, the share for software and databases shows a persistent upward trend, which in the US has surged since the late 1990s. In the UK, the rise in the software share was preceded by a contraction in hardware shares during the early 2000s, resulting in a more modest net ICT share relatively to the US. The key evidence is that while hardware is falling in all countries, software is rising possibly signaling a larger role for AI in the economy.

Figure 5 – Capital income shares (ICT, Hardware, Software) - US, UK and EU

Non-farm business sector



Source: author calculations from EUKLEMS and INTANProd database.

To investigate the impact of AI more deeply, we also analyse these dynamics within the industries responsible for producing AI technologies—the software-producing sector[2].

Figure 6 shows the income shares of all capital assets. There is a sharp fall and rise in the US software sector in the early 2000s, likely reflecting temporary changes in the rate of return over the "dot com" bubble burst period, since when shares have risen. The EU and UK software sectors have relatively lower capital income shares over the whole-time span.

Figure 7 shows the share of payments for hardware, software and R&D in the software-producing sectors. The US experienced the largest increase, with its total ICT share rising from 2001 onwards, driven by software and databases, while hardware remained flat. In the EU, the share stayed stable at approximately 0.10, showing little change over the two decades. Conversely, the UK saw a decline, with its ICT share falling from 0.10 to about 0.06.

The main takeaway from these data regarding levels is that in the three economies, software and databases have a larger capital payments share than physical IT equipment. Regarding trends, hardware capital income shares are falling, or flat in the software-producing sector. With the exception of the UK, software income shares are generally rising.

What does this mean for theories that rely on the elasticity of substitution between K and L? In a two input model the evolution of the capital payments shares provides information on this, but not necessarily in models with more than two inputs. Our finding here is that one has to be careful on the definition of payments to "computer" or "ICT" capital. If such payments are to hardware, that capital payments share is falling. If they are to software, it is rising.

,.

[2] The software producing sector for the US includes: NAICS Publishing industries, except internet (includes software) 5110; Data processing, internet publishing, and other information services 5140; Computer systems design and related services 5415. For the EU countries it refers to NACE Rev 2 sectors J58-J60 Publishing, audio-visual and broadcasting activities and J62-J63 IT and other information services.

Figure 6 – Capital income shares (Total assets) - US, UK and EU

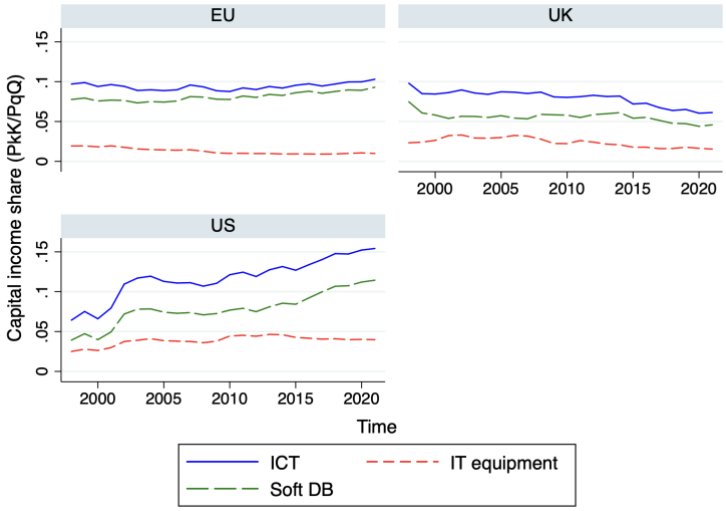
Software producing sector



Source: author calculations from EUKLEMS and INTANProd database.

Figure 7 – Capital income shares (ICT, Hardware, Software) - US, UK and EU

Software producing sector

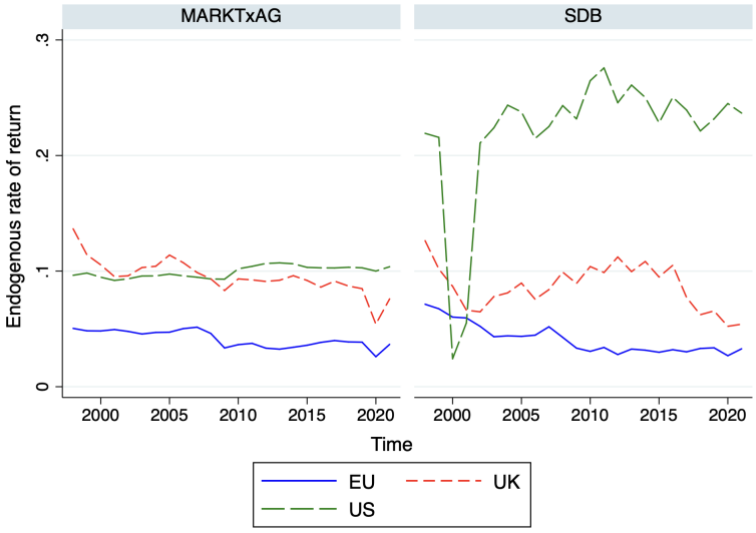


Source: author calculations from EUKLEMS and INTANProd database.

Figure 8 compares the Endogenous Rate of Return (ROR) across the US, UK, and EU, broken down by sector. Combined with the previous trends in capital income shares, the data reveals a significant difference in the rate of return between regions and asset types, specifically within the software sector.

Figure 8 – US, UK and EU, Total assets – Rate of return

Non-farm business vs software producing sector



Source: author calculations from EUKLEMS and INTANProd database.

While the return in the non-farm business sector remained relatively low and stable for all three regions—typically between 0.03 and 0.12—the software sector generated much higher returns in the US. Following a sharp drop during the dot-com bubble, the US rate of return for software recovered quickly and remained high, consistently staying above 0.20.

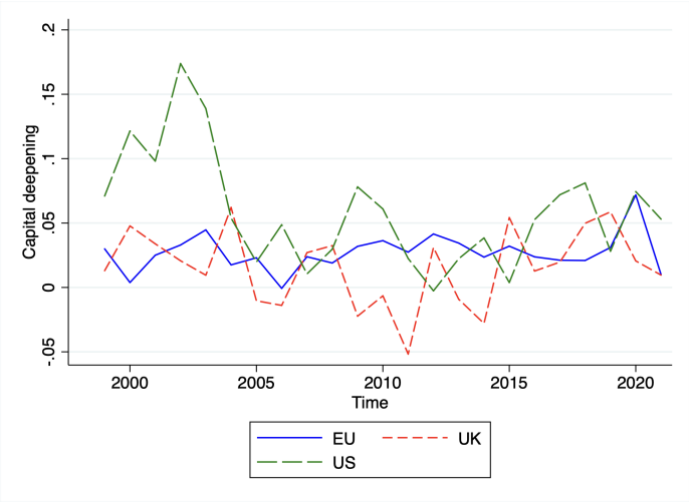
On the other hand, the UK and the EU economies did not show similar dynamics in the software sector; in the UK the ROR fluctuated at a lower level before declining after 2017, while in the EU consistently the ROR remained very low, eventually falling toward 0.03. These trends indicate that the US has been more effective at generating high returns from software and data, whereas in the EU and UK, these returns have stayed close to the levels of traditional assets.

6. Capital deepening in Software-Producing Industries: A Comparative Analysis

Finally, we look at the dynamics of hardware and software capital deepening within software-producing industries to get an initial indication about the different stages of development of the basic infrastructure required to produce AI across US, UK and EU (Figure 9 and 10). Data suggest different investment behaviours that likely reflect the varying maturity of the national infrastructure needed for AI production. Furthermore, the data reveal different investment

cycles across countries, as hardware and software capital deepening fluctuations are shifted over time. The United States shows strong hardware and software capital deepening in the early 2000s and a renewed upward trend after 2015. This expansion aligns with the growing infrastructure requirements for large-scale cloud computing, which serves as the foundational architecture for contemporary AI development.

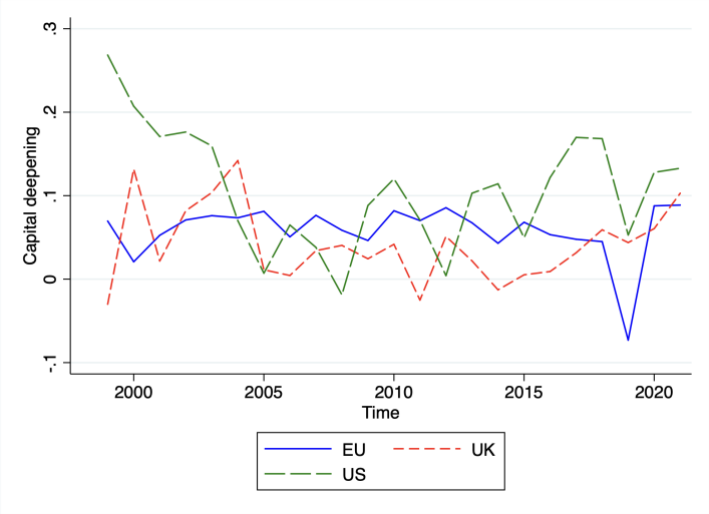
Figure 9 – Hardware Capital deepening – Software producing sector



Source: author calculations from EUKLEMS and INTANProd database.

Instead the EU countries maintain a stable and moderate trend, with both hardware and software capital deepening remaining relatively low; this suggests a relatively more gradual approach to AI related technological advancements. Finally, the UK shows greater volatility, characterized by periods of negative hardware capital deepening in which the physical asset stock contracted.

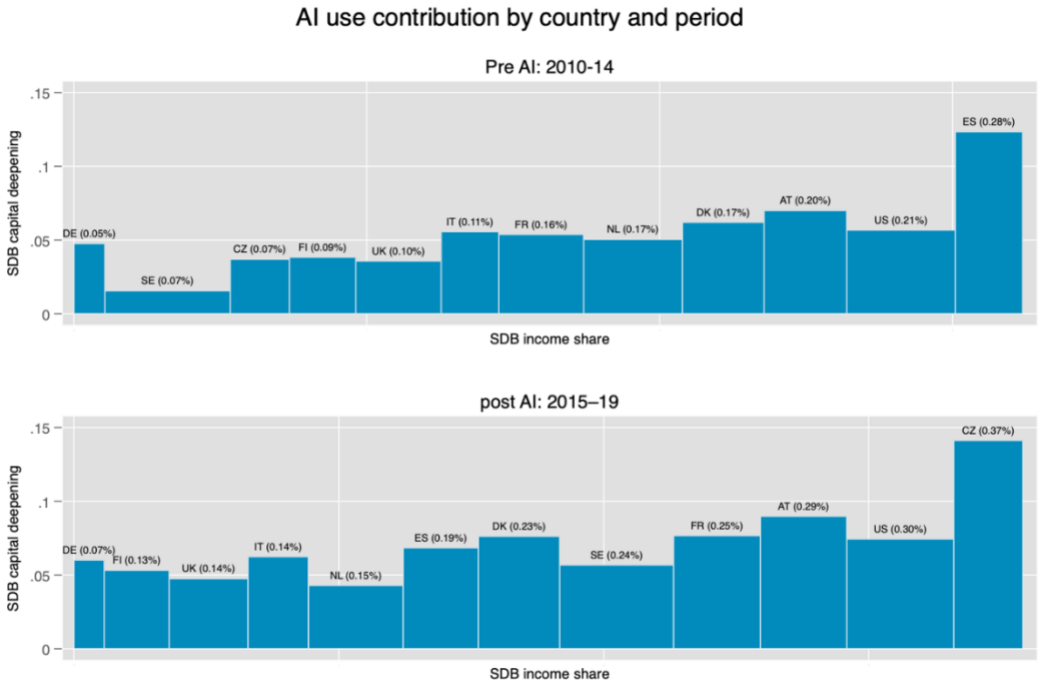
Figure 10 – Software Capital deepening – Software producing sector



Source: author calculations from EUKLEMS and INTANProd database.

We show some more detail for individual countries in Figure 11 which shows software capital deepening in the software producing sector for 2010-4 and 2015-19. We think of these as a pre- and post-AI period: we can look at more recent periods when we have more updated data. The width of each bar is proportional to the share of that sector in the country value added, and the height is that annual average growth in the software capital to hours ratio. The numbers at the top of each bar is the contribution to labour productivity growth, namely the share times the growth.

Figure 11 – Software Capital deepening – Software producing sector, by country



Source: author calculations from EUKLEMS and INTANProd database.

Almost every country increased the use of software and AI tools in the 2015–2019 period relative to previous years resulting in increasing growth contribution of the software sector.

7. Relative Prices and Financial Implications: Is AI a bubble?

The results described above suggest that the software-producing sector is (a) significantly more capital intensive than the software-using sectors, and (b) enjoys much higher ex post returns to its capital. Furthermore, we see indications that the US producing sector has become more even more capital intensive in recent years, and that, again in the US, there are early indications of a rise in software capital use in recent years, in part via the conduct of software/AI R&D in the using industry sectors.

These findings raise additional lines of inquiry centering on the financial implications of declining (user) software capital costs to be more fully explored in the final version of this paper.

7.1 A two-sector open-economy model

We will start by reviewing the economic fundamentals relevant to financial markets:

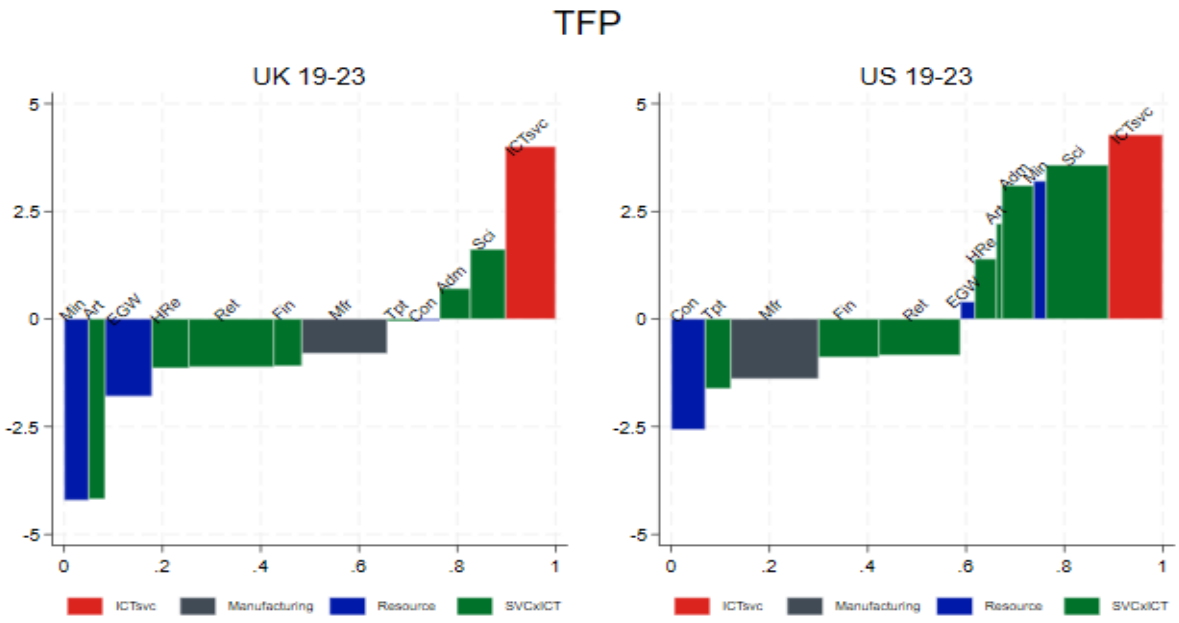
- Using a two-sector model, where a software/AI sector produces AI agents (software), while the rest of the economy uses the agents along with labor and tangible capital to generate consumption, a key element is the relative price of software to consumption goods, defined as $\mathbf{p} = \mathbf{P_S/P_C}$.
- The model posits that productivity growth is faster in the software sector, leading to a decline in the relative price of software capital goods over time. Profit maximization implies that the **real user cost of software capital** in the consumption sector equals the real interest rate plus depreciation minus the rate of relative price change. Thus, a falling **p** reduces the user cost of software capital, encouraging firms to accumulate more software capital and increase capital deepening. This process raises labor-productivity growth—and does so even in countries that do not produce (or design) the AI agents themselves.
- From an economic perspective, then, as software capital becomes more affordable relative to consumption goods, firms can allocate greater resources toward productivity-enhancing software. The resulting increases in productivity support higher earnings and bolster long-term equity returns.

Before discussing financial markets per se, we make another observation relevant to the broader economies of the US and UK (areas where we have more recent data as of this writing). A concern among financial analysts (as reported in the press) is that stock prices are elevated by outsized expectations for the return to AI, i.e.., that investments in AI are creating an asset price "bubble."

- Our earlier research estimated that roughly half of the growth in US labor productivity stemmed from the production and adoption of software capital. The finding does not imply, however, that the relatively strong recent growth of total factor productivity in the United States relies only on the digital sector.

  Figure 12 – Contributions TFP growth by Industry Sector

  in the UK and US.



Discuss this figure where we shall make the following points

- Chart shows industry-level TFP growth, where the width of each sector reflects the sector's production share
- Shows that the digital sector (red bar) is a key driver of TFP growth in both the US and UK.
- … note, however, that US economy is **not** entirely dependent on its digital economy (note the relatively tall and wide bars indicating strongly positive contributions from four other services sectors). In the UK, only its digital sector makes a strongly positive contribution to the overall change in TFP.
- The large finance, retail, and manufacturing sectors are negative contributors to TFP growth in both economies
- … but these 3 sectors are among **the 6 leading sectoral adopters of AI software** (according to IDC's AI software spending data).

7.2 Financial market implications

Are stock market valuations inflated due to an AI "bubble." We will contrast the economic fundamentals discussed above—that justify elevated valuations—with financial-market dynamics that suggest such valuations may be fragile.

- **Asset pricing: heterogeneity in risk and return.**

Investment-specific technology (IST) shocks—captured by declines in the quality-adjusted price of new capital goods—affect firms differently. Christiano and Fisher (2003) find that stock prices (the marginal cost of installed capital) are pro-cyclical while investment-goods prices are counter-cyclical; stock prices respond mainly to neutral technology shocks. Kogan and Papanikolaou (2012) show that positive IST shocks increase the market value of high-growth firms far more than low-growth firms. These growth-oriented firms earn different risk premia and tend to comove, giving rise to the value/growth factor. Applied to AI, this means the valuations of companies with large AI-enabled growth opportunities (e.g., cloud providers or AI-driven platforms) may legitimately exceed those of mature, capital-heavy firms, because their future cash flows are more sensitive to falling AI costs.

b. Financial-market dynamics: rational but fragile speculative growth.

Recent work by Caballero (2026) argues that high AI-related valuations can be part of a speculative-growth equilibrium: elevated valuations support rapid capital accumulation in AI, but they persist only as long as investors' optimistic beliefs are coordinated. AI capital behaves like labor, so it expands effective labor and keeps the marginal product of capital from falling. As wealth becomes concentrated among those benefiting from AI, aggregate saving rises and the required return falls, creating a funding feedback that sustains high valuations. However, these equilibria are fragile: if confidence falters, investment can collapse, leading to sharp price corrections. Thus, AI valuations may oscillate between optimism and pessimism even when long-term productivity gains are real.

c. Practical reading of market valuations.

Market commentary reinforces this nuance. Analysts note that AI-driven price/earnings ratios are roughly in line with long-term averages, suggesting current valuations are supported by profits and not wildly inflated.[3] The rapid rise in AI-related stocks reflects both hype and genuine earnings growth: The leaders are delivering profits, unlike many dot-com-era firms of the earlier IT era, but comparisons to past bubbles nonetheless highlight that sentiment can overshoot and then correct.

Accordingly, if artificial intelligence is indeed a transformative general-purpose technology underpinned by robust economic fundamentals, substantial capital allocation toward its development appears warranted. Nevertheless, it is important to recognize that financial market valuations associated with AI may exhibit considerable fragility. These valuations are inherently sensitive to shifts in investor sentiment, and the current trajectory of declining AI

[3] https://www.wellington.com/en/insights/chart-in-focus-is-ai-a-bubble-or-is-it-driving-real-market-value#:~:text=The%20rapid%20rise%20of%20AI,of%20innovation%20but%20also%20profit

capital costs—reminiscent of the reduction in technology prices during the ICT era—raises the prospect of speculative excesses and the potential emergence of an asset price bubble.

8. Declining Capital Costs and Measurement Challenges

All too often key economic measurements such as real GDP and productivity lag technological developments (Coyle 2025). The significance of relative price dynamics in reducing capital costs—and thereby driving AI-enabled gains in productivity and earnings—substantially heightens the importance of accurate price index construction.

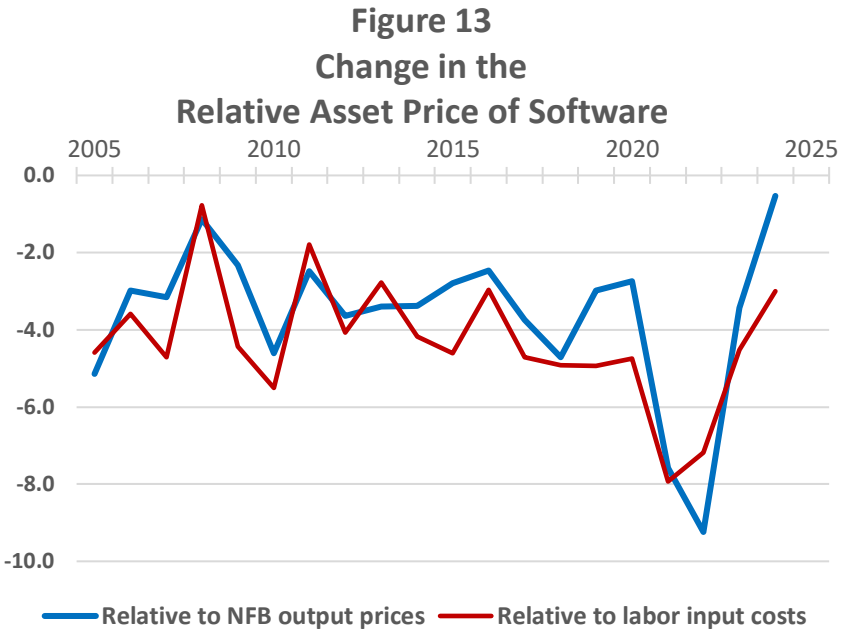
- Note: software/AI capital assets are of more than one type. Without regard to application area, we have
    - Legacy software
    - AI-enhanced software tools
    - Large-scale AI systems
- To our knowledge, price measures for large-scale AI system that dominate discussions of the transformation potential of AI do not exist; nor, to our knowledge do measures for smaller AI applications. Whether price measures for legacy software products that incorporate AI enhancements have been adjusted to reflect their changing character is also unclear, though it is possible that the collection of additional characteristics that would enable price collectors to do so is under way.
- Note: regarding consumption prices, cost-savings due to AI-enhanced process improvements in competitive, using industries—one of the major ways economists view the impact of AI (i.e., via automation)—will be reflected in lower prices for consumers.
- [carefully summarize/discuss the other scenarios: the surgeon using the robotic arm, not so much cost-saving, but improves the quality of the surgery, etc.]
- Finally, note that we should be taking GDP-B seriously. AI likely has already accelerated the availability of "free goods" to consumers: ChatGPT is a leading example. If the value of free goods has in fact accelerated in recent years, and to the extent it is likely to grow further, we will need to expand the boundary of GDP (consumption) to fully capture the impact of AI.

7.1 Software capital asset prices

The following figure shows two relative software price measures. [2025 and a measure relative to the PCE price index will be added],

- There is a slight deceleration in the relative price of software capital in recent years, especially when viewed relative to labor input costs, measuring as composition-adjusted hourly compensation.

- The dynamics of the relative price in recent years is less driven by the software asset price but rather by the numeraire, reflecting the step-up and subsequent deceleration of prices and wages in the United States.
- That said, it is unlikely that the capabilities of large-scale AI systems are reflected in these data.



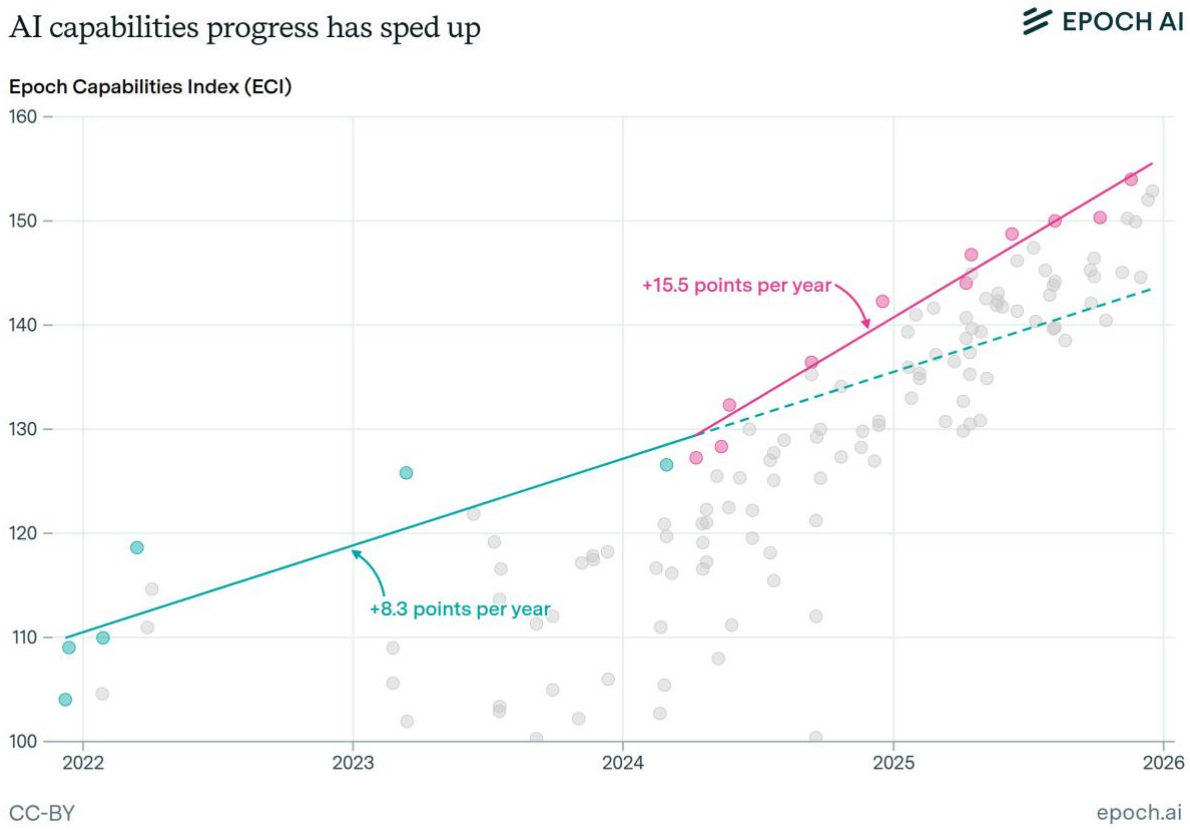
7.2 Characteristics of large-scale systems

The number of large-scale AI systems surged after 2022. Model-level performance characteristics, e.g., number of data points used, along with a host of related measures, are available from industry and academic groups dedicated to following frontier model development.

A model capabilities index that combines scores from many different AI benchmarks into a single "general capability" scale was developed late last year by Google researchers under contract to a non-profit research institute investigating the future of AI (EPOCH AI 2025; see the figure below). The measure is extremely helpful because it allows for comparisons between models over longer timespans than single benchmarks tend to reach saturation (about 3 years). The increase in the capabilities of leading-edge AI models, as reflected in the change in index points of the ECI, nearly doubled in 2025 after increasing 8.2 index points per year from 2022 to 2024.

The time span of the index also permits its use as quantity measure of large-scale AI systems. If $Rev^{LargeAI}$ represents the revenue attributable to the models in the EPOCH capabilities index

(ECI), then $P^{LargeAI} = Rev^{LargeAI}/_{ECI}$ represents a type of hedonic price index for large-scale AI systems.[4] We have not completed assembling the appropriate revenue stream to carry out this computation, but this can be done and will be analyzed in the final paper.



Source: Yafah Edelman and Jaeho Lee (2025), "AI capabilities progress has sped up". *Published online at epoch.ai.* Retrieved from: 'https://epoch.ai/data-insights/ai-capabilities-progress-has-sped-up' [online resource]

9. Preliminary conclusions

This outline has sketched a first pass at using the national accounts to track AI. We have argued that AI can be, in national accounts terms, be thought of as using a mix of software, hardware and databases to provide capital services and to build further AI. We have looked at trends in resulting shares in value added of payments to different capital assets across countries. For the nonfarm business economy as a whole, we have falling shares for hardware but rising

[4] Specially, a characteristic price index in which the quality-adjusted price is determined by a single characteristic, with hedonic function a ray emanating from the origin in characteristic space.

shares for software. For the software-producing sector, we have flat shares of hardware but, again, rising shares for software. Such changes are not driven by changes in the rate of return. Our preliminary conclusion is that, measured by capital income, software is eating the world but slowly.

We also outlined the implications of our analysis for interpreting investor returns and financial market valuations. The literature indicates that the valuations of AI model producers may be inherently fragile due to their sheer magnitude; nevertheless, such elevated valuations are supported by robust underlying economic fundamentals. At least in part, these fundamentals are manifest by an ongoing decline in the relative cost of software capital, specifically, software user costs measured in consumption units. Notably, the price measures employed in our national accounts do not seem to capture the enhanced capabilities of large-scale AI systems. Since 2022, the number of these advanced systems has accelerated, implying that the true extent of the reduction in relative software capital costs is understated in official statistics.

References

[to be supplied in the next draft]