Robert Solow once said, "You can see the computer age everywhere but in the productivity statistics.". Was he right, and will the rise of AI be much the same? What are the broader lessons we should draw for the theory and data-analysis of technical change?

Introduction:

Solow (1987) made his famous statement, "[Y]ou can see the computer age everywhere but in the productivity statistics," in his review of Cohen and Zysman's work on the post-industrial U.S. economy. His observation highlights a perceived discrepancy in the U.S. economy: despite massive investment in information technology (IT) since the 1970s, it failed to translate into positive labour productivity growth (Solow, 1987). As shown in Figures 1 and 2 (see Appendix), the exponential growth in real IT investments is associated with stagnant productivity growth, particularly in the service sector (Brynjolfsson and Yang, 1996).

Solow's observation sparked an academic debate about the relationship between IT investment and its impact on productivity growth. Scholars have conducted empirical studies, both supporting and refuting this observation. In economy-wide studies, empirical evidence is highly in support of Solow's observation. The work by Jorgenson and Stiroh (1995) substantiates that average multi-factor productivity growth fell from 1.7% per year in 1947-74 to about 0.5% in 1973-1992. Meanwhile, OCAM capital as a percentage of producers' durable equipment investment rose from about 0.5% in the 1960s to 12% in 1993. Stephen Roach (1987) compared the productivity of information workers across industries with that of production workers. He showed that productivity in these sectors not only failed to improve but actually declined, with output per information worker dropping by 6.6%.

However, a key limitation of these studies is their high level of aggregation, as other variables beyond IT investment could also influence labour productivity at the national level, leading to biased results. In contrast, industry-level studies are less prone to bias, and empirical evidence at this level yields diverging results. Productivity stagnation is mostly concentrated in the service sector (Schneider, 1987). Roach (1987, 1991) argued that IT reduced manufacturing labour but increased white-collar jobs in services. Similarly, Berndt and Morrison (1995) found no significant difference in the productivity of IT capital compared to other types of capital in most of the 20 industry categories they examined. However, they did find that IT investment was strongly correlated with a significantly increased demand for skilled labour.

How can the divergent findings between economy-wide and industry-level studies be reconciled? Is the so-called "productivity paradox" a genuine contradiction between economic theory and empirical evidence, or does it simply reflect broader economic dynamics? This essay argues that the productivity paradox is more of an observation than a true paradox. While technical change can increase productivity, it is offset by the rise of low-productivity jobs, zero-sum competitive activities, and changes in consumption patterns. This offsetting dynamic applies to all General Purpose Technologies (GPT), including the recent emergence of artificial intelligence (AI). The essay is organised as follows: Section I argues against the validity of the productivity paradox and evaluates the impact of the rise of AI on labour productivity growth. Section II examines lessons learnt from theories and data analysis of technical change, and Section III concludes the essay.

Section I: The productivity paradox and the rise of AI

To reconcile the discrepant findings from economy-wide studies with industry-level studies, one must consider the relative shares of these sectors in the total economy. Automation in one sector of the economy could free up labour and shift it to other sectors, and total productivity growth is as much driven by the productivity and productivity growth potential of the sectors into which workers move, as in the sectors where jobs are automated away (Turner, 2018). In this case, Baumol's cost disease helps explain why the relative share of sectors with rapid productivity growth and highly automatable jobs is likely to shrink over time, while sectors with slower productivity growth and less automatable jobs will absorb more labour and expand their relative size. In a multi-sector economy, different sectors experience productivity growth at varying rates (Baumol, 1967). Some sectors, such as manufacturing and certain automated services, are highly amenable to productivity gains through automation and technological innovation. Others - particularly those reliant on human interaction, creativity, or tasks that require a high level of finesse, such as healthcare and education - are less susceptible to automation and thus experience slower productivity growth.

As automation enhances productivity in capital-intensive sectors, fewer workers are needed to produce the same output. This displacement of labour leads to a sectoral shift, where workers who lose their jobs in high-productivity industries must seek employment elsewhere (Turner, 2018). Due to skill mismatches, a majority of displaced workers often end up in low-productivity, low-paid jobs that are labour-intensive and resistant to automation, such as custodial work, food service, and delivery. In these sectors, productivity gains are difficult, if not impossible, to achieve because they rely on human labour for their core value proposition, which machines cannot easily replicate.

In addition to low-productivity, low-paid jobs, displaced workers and future labour force members increasingly shift into zero-sum activities - sectors where competition reallocates resources without generating new value, such as legal services, finance, advertising, and marketing (Turner, 2018). In these sectors, the economic activity primarily revolves around shifting or protecting economic gains rather than creating new wealth. For instance, corporate lawyers may engage in protracted litigation battles over intellectual property, while financial traders speculate on assets, none of which directly increases the production of goods or services that improve human welfare. These zero-sum activities proliferate because they are relatively insulated from automation; the human judgment and strategic thinking required in such fields are difficult to replicate with machines.

Moreover, demand for zero-sum activities is highly income elastic and price inelastic (Baumol, 1967). As incomes rise, wealthier individuals allocate more of their income to these services, which are often seen as luxuries, driving demand faster than income growth. These services are also viewed as essential or status-driven, making them price inelastic; even as prices increase, demand remains steady or grows due to their perceived necessity or exclusivity.

Collectively, the displacement of labour into low-productivity sectors, the growing attractiveness of zero-sum activities, and shifts in consumption patterns drive the economy toward asymptotic productivity growth. As high-productivity sectors become more capital-intensive and less reliant on human labour, their contribution to overall employment diminishes. Meanwhile, low-productivity sectors expand, but without the same potential for productivity gains. This creates a scenario where even though technological innovations continue to enhance efficiency in certain parts of the economy, the overall growth in productivity slows. Essentially, as more workers are absorbed into low-productivity or zero-sum jobs, the economy's aggregate productivity stagnates, asymptoting towards zero growth over time.

The development of AI follows a similar trend. Figure 3 illustrates the annual percentage change in labour productivity for the U.S. non-farm business sector from 2009 to 2024. After 2009, productivity growth declined rapidly and stayed low, stagnating near 0% from 2011 to 2019. The COVID-19 pandemic in 2020 brought extreme volatility, leading to a sharp collapse in 2022 and a major decline in productivity gains.

Academic debates on the causes of the current decline in labour productivity with the rise of AI are divided into two opposing arguments. Technological optimists focus on theoretical models that highlight the long-term potential of AI to drive productivity gains. Scholars like Aghion et al. (2017) and Acemoglu & Restrepo (2018) suggest that AI could eventually lead to significant productivity increases, but they acknowledge that these benefits may take time to materialise due to implementation lags and the need for complementary innovations. They argue that AI's transformative power could even lead to a technological "singularity", where rapid economic growth ensues once these technologies are fully integrated into the economy. In contrast, Tech pessimists argue that despite AI's advancements, its impact on productivity remains limited, aligning with the broader theory of secular stagnation, as proposed by Gordon (2016). They contend that deeper structural issues, such as demographic changes, reduced capital investment, and weak demand, are keeping productivity growth low, and AI alone is insufficient to reverse this trend.

Aside from the aforementioned explanations for asymptotic productivity growth, the problem of mismeasurement and subsequent time lags may explain why AI's transformative potential has yet to clearly impact traditional productivity statistics. According to Brynjolfsson, Rock, and Syverson (2017) , a key challenge in measuring AI's impact is that intangible capital - such as AI technologies, data infrastructure, and new organizational processes - often goes unmeasured or is inadequately captured in productivity statistics. Traditional measures focus primarily on tangible capital like machinery and labour, thus failing to account for the contributions of these intangible assets that are essential in knowledge-driven economies. The rise of AI exacerbates this issue, as investments in AI require complementary adjustments which take time to mature and may not immediately reflect in productivity numbers. The mismeasurement problem is also reflected in the Solow Residual, a standard measure of productivity growth that captures the portion of output growth not explained by changes in capital or labour inputs and serves as a proxy for technological progress (Brynjolfsson et al., 2017). When intangible capital like AI is not measured, the Solow Residual can misrepresent actual productivity trends, particularly in the early stages of AI adoption. During this particular period, AI investments are increasing, but the benefits have not yet materialised, leading to an underestimation of technological progress. As AI investments begin to deliver returns, however, the Solow Residual may reflect an overestimation of productivity growth, as the delayed benefits from earlier investments are finally realized.

Additionally, the time lag associated with AI adoption plays a critical role. As a general-purpose technology (GPT), AI requires a prolonged adjustment period before its full potential is realised. Firms and industries must undergo extensive restructuring, such as changing business models and training employees, processes that can take years or even decades. This slow adjustment phase is represented by the J-Curve effect (figure 4), where the high costs of implementing AI - like retraining workers - can lead to an initial decline or stagnation in productivity, forming the downward slope of the J-Curve. Over time, as AI becomes fully integrated and complementary innovations take hold, productivity gains accumulate, leading to an upward trajectory in productivity growth, marking the rising slope of the J-Curve and eventually resulting in net productivity increases.

Section II: Implications from theories and data analysis of technical change

There were 3 major technological changes in economic history, the First Industrial Revolution, driven by the mechanisation of production and the rise of steam power in the late 18th century, the Second Industrial Revolution, characterized by the expansion of electricity, mass production, and railroads in the late 19th century, and the Third Industrial Revolution, marked by the digital revolution and the rise of information technology in the late 20th century. Theories and data analysis regard these consistently demonstrate a cyclical pattern of technological diffusion, which unfolds in three distinct stages: recognition and introduction, production synergy, and maturity. In the initial stage, businesses identify the potential of new General Purpose Technologies (GPTs) but adoption is slow due to high uncertainty and substantial fixed costs, limiting immediate productivity gains. As diffusion enters the production synergy phase, economies of scale begin to materialise, reducing marginal costs, and driving higher rates of investment in complementary innovations, infrastructure, and human capital. This stage also witnesses an acceleration in TFP growth as industries optimize the use of the new technology, both within core sectors and as it spills over into peripheral industries, fostering sectoral linkages that enhance aggregate economic output. However, diffusion is not uniform; network externalities and skill-biased technological change can lead to inequality, with certain traditional or capital-intensive sectors lagging behind (Turner, 2018). The final maturity stage is marked by diminishing returns on further technological adoption as the technology saturates the market, resulting in a deceleration of productivity growth. Nonetheless, these GPT-driven cycles typically lead to structural shifts in the economy that reallocates labour and capital and gives rise to entirely new industries while obsolescing others.

Section III: Conclusion

The essay evaluates the so-called 'productivity paradox' and reveals that it is more of an observation than a true contradiction between theory and empirical evidence. The paradox stems from structural shifts in the economy, where technological innovations boost productivity in certain sectors but displace labour into the service sector, where low-paid, low-productivity jobs and zero-sum activities proliferate. The rise of AI follows a similar trajectory: it holds transformative potential for future productivity growth but this growth is constrained by sectoral imbalances, time lags in adoption, and mismeasurement issues. For all major technical change in history, all them followed a cyclical diffusion pattern with three stages: recognition, production synergy, and maturity.

Word counts: 2,076

APPENDIX:

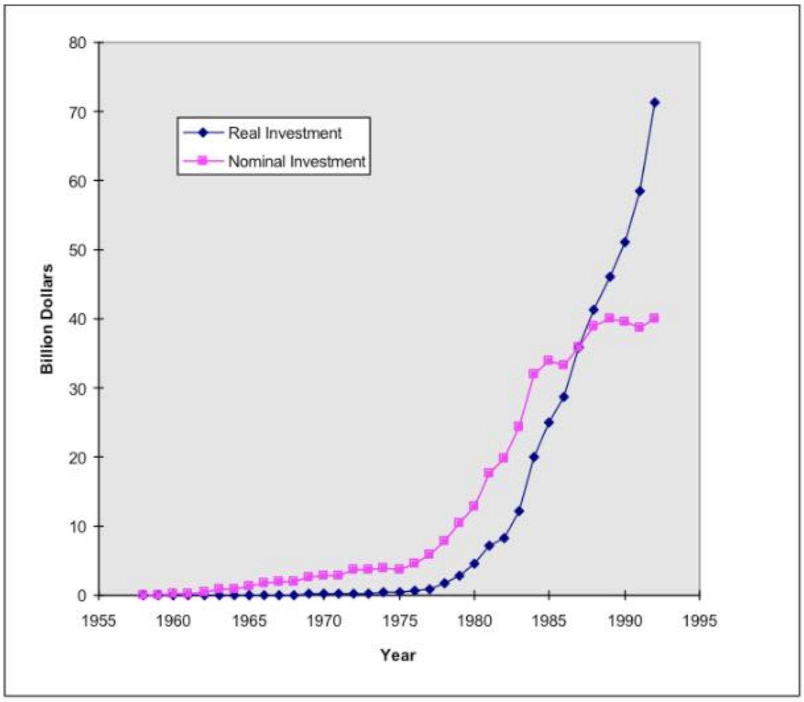

Source: Based on data from [BEA, National Income and Wealth Division], adapted from Jorgenson and Stiroh [1995].
Note: Constant dollars (base year 1987) calculated by hedonic price method, see Dulberger [1989].

Figure 1: Investment in Information Technology (measured in billions of dollars) from 1955 to 1995.

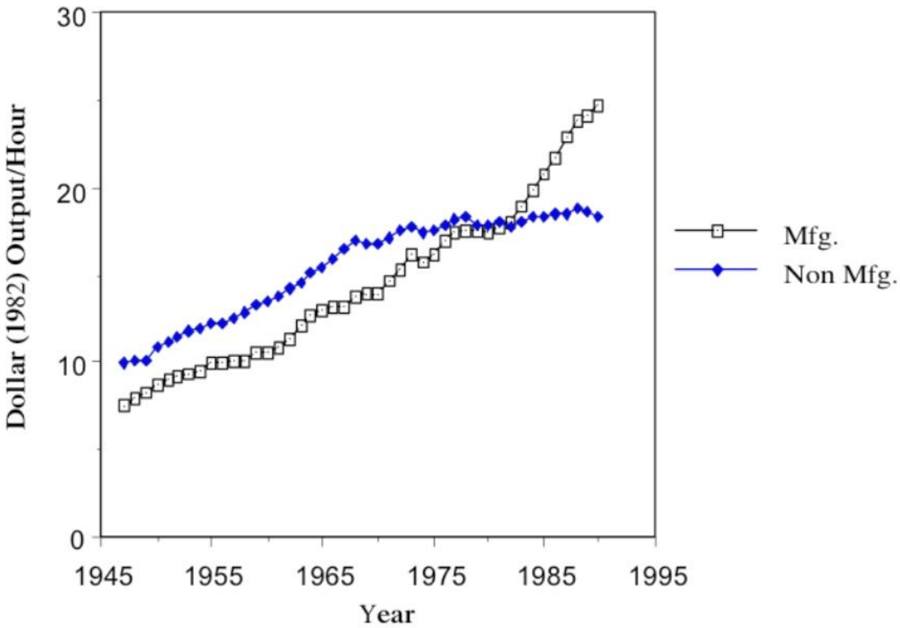

Source: Based on data from [Bureau of Labor Statistics, Productivity & Testing]

Figure 2: labour Productivity (measured in dollar output per hour) in the manufacturing sector and the service sector from 1945 to 1995.

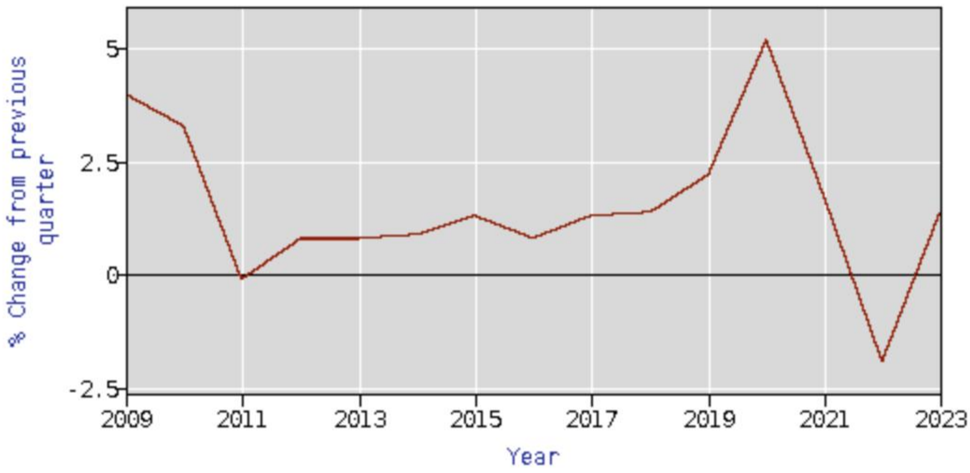

Figure 3: the annual percentage change in labour productivity for the U.S. non-farm business sector from 2009 to 2024.

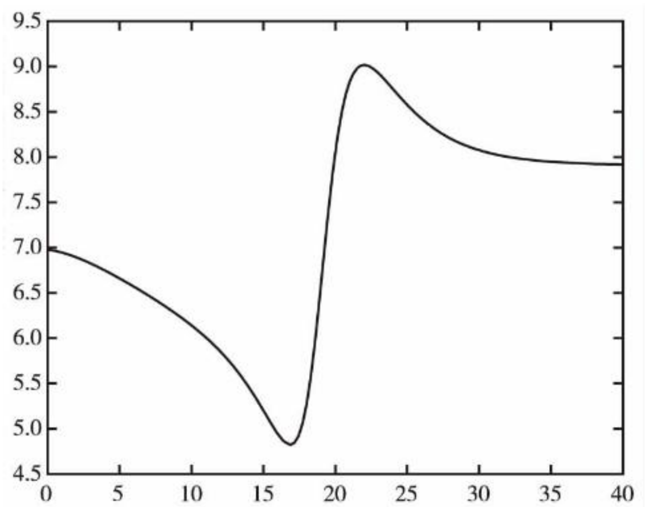

Figure 4 illustrates the J-curve effect, where the x-axis represents time in 5-year increments, and the y-axis represents labour productivity.

Bibliography:

Acemoglu, D. and Restrepo, P., 2018. Artificial intelligence, automation and work. *NBER Working Paper,* No. 24196.

Aghion, P., Jones, B. and Jones, C. (2017) Artificial Intelligence and economic growth.

Baumol L.W. (1967) 'Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis', *The American Economic Review,* Vol. 57, No. 3 (Jun., 1967), pp. 415-426

Berndt, E.R. and Morrison, C.J. (1995) 'High-tech capital formation and economic performance in U.S. Manufacturing Industries an exploratory analysis', *Journal of Econometrics*

Brynjolfsson, E. and Yang, S. (1996) 'Information Technology and productivity: A review of the literature', *Advances in Computers,* pp. 179-214.

Brynjolfsson, E., Rock, D. and Syverson, C. (2017) Artificial Intelligence and the modern productivity paradox: A clash of expectations and statistics.

Gordon, R.J. (2016) *The rise and fall of American growth.* Princeton University Press.

Jorgenson, D.W. and Stiroh, K. (1995) 'Computers and growth', *Economics of Innovation* *and New Technology,* 3(3-4), pp. 295–316.

Roach, S.S., 1987. ‘America’s technology dilemma: a profile of the information economy.’ *Morgan Stanley.*

Schneider, K., 1987. ‘Services hurt by technology: productivity is declining.’ *The New York Times,* 29 June, pp. D1, D6.

Solow, R.M. (1987) ‘We’d better watch out,’ *the New York Times Book Review,* 36, p. 36.

Turner, A., 2018. Capitalism in the age of robots: work, income and wealth in the 21st century.

Roach, S.S., 1991. ‘Services under siege: the restructuring imperative.’ *Harvard Business Review,* 39(2), pp.82-92.