The Nature of technology, structural transformation, and the nature of work

1 Introduction

Technology is a fickle thing. During the century after 1870, rapidly advancing, labour-saving technology pulled workers into the manufacturing sector. This didn't fundamentally change the nature of work – most people were still working with their hands – but it changed what they were making and where. During the same period, labour's share of GDP was stable, and income inequality was falling. And then the correlations flipped.

During the half century following 1970, rapidly advancing, labour-saving technology pushed workers out of manufacturing. It also deeply transformed the nature of work for many. Most workers switched to service-sector jobs that involve less 'hand work' and more 'head work'. At the same time, labour's share of GDP fell, and inequality rose (see the charts).

How could technology have such radically different effects during these two different eras?

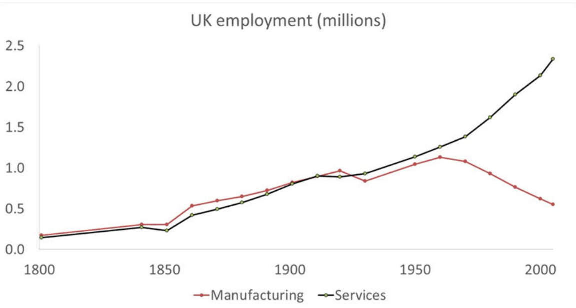

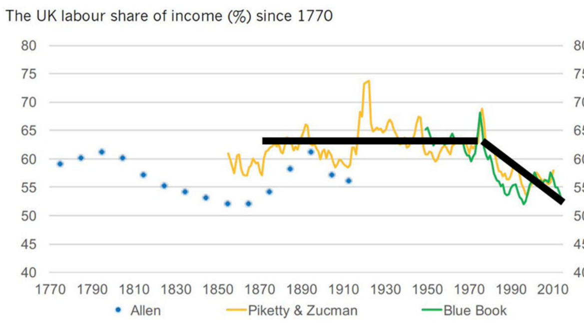

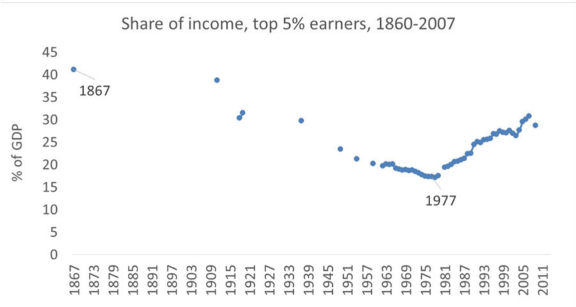

This paper, which is very much in draft form, has two goals. The first is to crystallise – in a parsimonious general equilibrium setting – a simple accounting of the differences between the jobs-technology link from 1870-1970, and the link from 1970-present. The account rests on the assertion that that outcomes were radically different because the technological shocks were radically different. The pre-1970 developments were driven by mechanisation; the post-1970 outcomes were driven by computerisation.

Mechanisation favoured the making of things. The resulting automation directly boosted the productivity of people who worked with their hands. It helped people who worked with their heads, but only indirectly because it was a technology of things, not thoughts. It created masses of new industrial jobs. Moreover, since most people back then worked with their hands, the more-manual-than-mental aspect of the mechanisation lowered inequality by boosting the relative productivity of 'hand workers' versus 'head workers'. The 1970s technological impulse did just the opposite.

Computerisation provided better tools for those who work with their heads, but better replacements for those who work with their hands. Creating better replacements for factory workers – robots and the like – was a massive push-factor that emptied factories even as industrial output continued to rise. The better tools for brain workers, by contrast, was a massive pull-factor into the service sector. It created millions of new service-sector and professional jobs – many of them in occupations that were previously unimageable. From the social cohesion point of view, the new technology was divisive. Since the 'head workers' were already better off than the 'hand workers', a technology which favoured brain over brawn favoured the few who were already favoured, while disfavouring the many who weren't.

The second goal is to contribute to the 'future of work' discussion by using general-equilibrium insights to address how the nature of future technology will affect the nature of future work. While we cannot know the names of future jobs, general equilibrium considerations allow us to think about the nature of those jobs. Here again, the main point of departure is the assertion is future technology is like to be very different than computerisation, or mechanisation.

The recent commercialisation of Artificial Intelligence (AI) – most notably in the form of machine learning – is bringing a radically different automation to the workplace. In a nutshell, AI-trained computer models are allowing 'software robots' to 'think' in ways they never could before. Machine learning combine with big data has given computers a brand-new range of cognitive skills and some of these are useful in workplaces, especially offices. Machine learning, in other words, is providing better substitutes for 'head workers' – or at least when it comes to certain tasks that previously could only be done by humans. This could have big effects for a very simple reason. 80-90% of the workforce is in the service sector.

Automation of factories will surely continue, but since such a small share of the 21st century workforce is in factories, the impact on office work is what will really affect the nature of work for societies as a whole. The key point is that many 'thinking tasks' are being automated – a range of tasks that used to require humans can now done by 'white collar robots' such as Robotic Process Automation systems, virtual assistants, and sophisticated AI platforms like IBM's Watson. This wave of technology is, in other words, is providing better substitutes for people who worked with their heads because it was a technology of cognition, not construction. This is likely to bring about important changes.

Until recently, most white-collar, service-sector, and professional jobs were shielded from automation by humans' cogitative monopoly. Computers couldn't think, so jobs that required any type of thinking – be it teaching nuclear physics, arranging flowers, or anything in between – required a human. Automation was a threat to people who did things with their hands, not their heads. Indeed, in many empirical studies, routine-ness is still used as a proxy for automation. Machine learning is changing this.

Plan of paper. The paper is organised in three substantive sections in addition to the introduction and concluding remarks. The first section looks at implications of mechanisation, the second looks at computerisation, and the third uses the same framework to look at the potential impact of AI.

Literature review (incomplete). The first two sections could be viewed as working in the structural transformation literature. Here some key papers are including Matsuyama (1992, 2009), Hémous, and Olsen (2014), Caselli and Coleman (2001), Ngai and Pissarides (2007), Herrendorf, Schmitz, and Teixeira (2012), Michaels, Rauch, and Redding (2012), , Bustos, Garber, and Ponticelli (2017), McMillan, Rodrik, and Sepulveda (2017), Eckert and Peters (2018), Karádi and Koren (2018), and Sotelo (2018). See Herrendorf, Rogerson and Valentinyi (2014) for an overview.

Matsuyama (2009) for a very succinct modelling of the two traditional drivers of structural change that Acemoglu and Guerrieri (2008) call "demand-side" (non-homothetic preferences) and "supply-side" (biased technical progress) explanations. Ngai and Pissarides (2007) refer to them as the "utility-based" and "technological" explanations, respectively.

In terms of modelling of the technological differences between mechanisation and computerisation, the key references are various papers by Acemoglu and co-authors, especially Acemoglu and Restrepo (2016) who distinguish 'enabling technologies' (factor augmentation) and 'replacing technologies' (factor replacing). Acemoglu and Autor (2010) provide a clear, and empirically motived discussion of the discussion of the differences and economic mechanism at play.

These articles build out from the very extensive empirical literature on the Skill-Biased Technical Change (SBTC) hypothesis posited by, inter alia, Freeman (1976), Katz and Murphy (1992), and Autor et al. (1998, 2008), and tested empirically by, for example, Katz and Murphy (1992), Autor Katz and Krueger (1998), and Autor, Katz and Kearney (2008), Hemous and Olsen (2014). The hypothesis is that a burst of new technology caused a rise in the demand for highly skilled workers that outstripped the supply and thus created inequality.

A related set of contributions in the trade theory literature concerns the impact of technology on wages and outputs. This work last garnered substantial attention during the late 20th century when the issue was whether the deleterious developments in rich nation labour markets were due to 'skill biased technological progress' or globalisation. Most notable was the robust Krugman-Leamer exchange of views, e.g. Leamer (1994, 1996) and Krugman (1995). Other noteworthy contributions were made by Wood (1994), Jones and Ruffin (2007), and Xu (2000).

2 Technology as machines helping manual workers

A worker with an electric drill can do the work of, say, three workers using hand drills. At least since Robert Solow's famous work on growth in the 1950s, economists tend to think of technological progress in this way – machines helping humans do more (this is also called the 'factor augmenting' view).

From a partial equilibrium perspective, this sort of innovation cuts both ways when it comes to jobs (Ngai and Pissarides 2007). Mechanization means that the same pile of work can be done with fewer workers. All else equal, that makes technology into a job-destroying force. But the implied cost savings means lower prices and thus more sales (a higher pile of work) and this tends to make technology into a job-creating force. There is, in a sense, a footrace between the height of the pile and efficiency of workers. When the productivity-production footrace is won by the piling-raising side – technology acted as a 'pull factor' – pulling workers into the sector. Where the efficiency-side wins, it is a 'push factor', pushing workers out of the sector.

This partial equilibrium view misses some critical elements. If technology is a push factor, where do the displaced workers go? If it's a pull factor, where do they come from? This is where general-equilibrium (GE) analysis becomes important. Standard GE insights, however, are difficult to apply since GE insights link price and endowment shocks to wages, and outputs – not jobs. There is, however, an intellectual sleight-of-hand – first demonstrated by Jones (1965) – that gets past this.

A few moments reflection reveals that under the machines-helping-humans conceptualisation of progress, one worker with a better drill can be thought of as three workers using the old technology. This reflection – the Jones' Equivalence –reduces a new problem – what will better robots do to output? – to a problem previously solved – what happens to output when factor endowments change? The solution to the 'problem previously solved' has a name – the Rybczynski theorem. It says that an increase in the supply of one factor tends to expand the output of the sector that in intensive in its use. Moreover, the expansion is more than proportional, and other sectors tend to contract.

A slight twist on this result gives us the jobs impact. As it turns out, the employment impact turns on the magnification aspect of Rybczynski. A tech shock that is Jones-equivalent to an expansion of the supply of A-workers will expand output of the A-intensive sector more than proportionally. This means it will create new jobs for A-workers in the A-intense sector since the output expansion more than compensates for the efficiency gain. This provides insights about the nature and sector of future jobs – even if it does not tell the names of the new jobs.

This GE insight may seem counterintuitive from the partial equilibrium angle. After all, it says that the new jobs, the jobs of the future, will be relative intensive in the use of the type of labour that robots are replacing today. And the new jobs will appear in sectors that are relatively intensive in the type of labour with which robots are competing most directly. The insight, by contrast, is obvious from the general equilibrium perspective. If an economy gets more unskilled workers, the unskilled-intensive sector must expand to absorb them all. Since this expansion will require some skilled workers – and these need to come from somewhere – the skill-intensive sector must shrink.

This section presents a minimalist framework that permits crystallisation of the general equilibrium insights. It offers an economically logic accounting for why jobs were created fastest for workers who were most directly concerned with the progress and in sectors where these workers were used intensively. In short, it is an accounting of how mechanisation led to an expansion of the manufacturing sector. The model presents a version of the Rybczynski analysis in Baldwin, Haaland and Venables (2019); it is simplified in terms of the technology but set in an imperfectly competitive setting that allows us to talk about new firms, new products, and new jobs.

2.1 The basic GE model – Helpman-Krugman meets Jones

The economy has two productive factors (L_A and L_B) and two sectors (1 and 2) both of which are marked by Dixit-Stiglitz monopolistic competition. Preferences are CES within sectors and Cobb-Douglas across sectors with μ as the expenditure share on sector-2 varieties.

Production of each variety is marked by increasing returns defined by homothetic, linear cost functions: $c_i[w_A, w_B, x_i] = (w_A a_{Ai} + w_B a_{Ai})(a_i x_i + F_i)$; $i = 1,2$. Here the w_j $(j = A, B)$ are wages, the a_{ji} $(j = A, B, i = 1,2)$ are the unit, labour-input requirements, and a_i and F_i are the variable- and fixed-costs parameters. Sector-1 production is intensive in 'manual' tasks and sector-2 is intensive in 'thinking' tasks; A-workers are abundant in manual 'skills' (M-skills) and B-workers are abundant in 'thinking' skills (T-skills). The factor-intensity assumption, written using the factor-intensity ratios, $\lambda_j \equiv (a_{Aj}/a_{Bj})$, is $\lambda_1 > \lambda_2$, so A-workers are used intensive in sector 1 while B-workers are used intensively in sector 2.

This section simplifies by assuming A-workers have only M-skills, and B-workers have only T-skills, so there is a perfect congruence between worker-type and sector skill-intensity. The regularity condition that ensures firms operate in both sectors in equilibrium is $\lambda_1 > \Lambda > \lambda_2$, where Λ is the endowment ratio.

As per Dixit-Stiglitz monopolistic competition, each firm produces a unique variety, and optimally prices it as a mark-up over marginal costs, i.e. $p_i = (w_A a_{Ai} + w_B a_{Ai})(a_i/(1-1/\sigma_i))$, where σ_i is the sector-i CES elasticity of substitution. Choosing units such that $a_i/(1-1/\sigma_i) = 1$, and solving the free entry conditions, the equilibrium scale of firms, $\bar{x}_i$, and prices, p_i, are $p_i = w_A a_{Ai} + w_B a_{Ai}$, $\bar{x}_i = F_i \sigma_i;\ i = 1,2$. With this, the closed economy equilibrium is characterised by employment, pricing, and market-clearing conditions, namely:[1]

$$\boldsymbol{p} = \boldsymbol{A}\boldsymbol{w}, \qquad \boldsymbol{L} = \boldsymbol{A}'\boldsymbol{X}, \qquad p_2 = \mu X_1/(1-\mu)X_2 \tag{2}$$

where the prime indicates the transpose and:

$$\boldsymbol{w} \equiv \begin{bmatrix} w_A \\ w_B \end{bmatrix}, \quad \boldsymbol{p} \equiv \begin{bmatrix} 1 \\ p_2 \end{bmatrix}, \quad \boldsymbol{L} \equiv \begin{bmatrix} L_A \\ L_B \end{bmatrix}, \quad \boldsymbol{X} \equiv \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}, \boldsymbol{A} \equiv \begin{bmatrix} a_{A1} & a_{B1} \\ a_{A2} & a_{B2} \end{bmatrix}, \ X_i \equiv n_i \bar{x}_i$$

Here n_i is the mass of products/firms in sector-i.

Clarity is served by adding a few assumptions that unclutter the formulas in this and subsequent sections by eliminating asymmetries that are unrelated to factor intensity. The number of A and B workers is equal ($L_A = L_B \equiv L$), expenditure shares are equal ($\mu = 1/2$), and the endowment ratio is such that $\Lambda = (\lambda_1 + \lambda_2)/2$. Solving for the number of firms/products and the number of jobs per sector per labour type yields:

$$n_1 = \frac{a_{B2}L_A - a_{A2}L_B}{a_{A1}a_{B2} - a_{A2}a_{B1}} \frac{1}{F_1\sigma_1}, \qquad n_2 = \frac{a_{A1}L_B - a_{B1}L_A}{a_{A1}a_{B2} - a_{A2}a_{B1}} \frac{1}{F_2\sigma_2}, \qquad J_{ij} = a_{ij} n_j F_j \sigma_j \tag{3}$$

where J_{ij} is the number of jobs for i-type workers in sector j.

Observe how these solutions for the n's, and thus output and jobs, are determined entirely by employment considerations. This narrows the focus to the analysis to the pure, GE effects; the n_i's, outputs, and thus jobs, respond to neither relative prices nor relative wages. As a result, the classic, structural-transformation considerations of demand elasticities and factor substitutability (Ngai and Pissarides 2007) are left aside. This is intentional and it is meant to clarify the GE mechanisms of action.

2.2 Conceptualising technology – machines helping humans do more

Technological progress is conceptualised as reductions in the unit labour-input requirements (the a_{ji}'s). The Jones' equivalences between such factor-augmenting tech shocks and endowment and

[1] All section 1 varieties always have the same price so there is no ambiguity as to the numeraire, $p_{1=1}$.

price shocks are simple.[2] Technical progress means a fall in some or all of the a_{ij}'s. The vector of labour savings, $\boldsymbol{S_Q}$, and vector of cost savings, $\boldsymbol{S_C}$, are:

$$\boldsymbol{S_Q} \equiv \begin{bmatrix} S_{QA} \\ S_{QB} \end{bmatrix} = \boldsymbol{A_S'} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}, \ \ \boldsymbol{S_C} \equiv \begin{bmatrix} S_{C1} \\ S_{C2} \end{bmatrix} = \boldsymbol{A_S w} \quad \text{where } \boldsymbol{A_S'} \equiv \begin{bmatrix} da_{A1} & da_{A2} \\ da_{B1} & da_{B2} \end{bmatrix} \tag{4}$$

and all elements of the $\boldsymbol{A}$ and the $\boldsymbol{S}$'s are positive.

With these, the employment condition becomes $\boldsymbol{L} + \boldsymbol{S_Q} = \boldsymbol{A_S}'\boldsymbol{X}$, so we can view any tech shock as an endowment shock, and thus use the Rybczynski insights to address the jobs question. The pricing equations become: $\boldsymbol{p} + \boldsymbol{S_C} = \boldsymbol{Aw}$, so the Stolper-Samuelson insight speak directly to tech-wages questions.

Direct calculation using (3) and (4) yields:

Result 1: Rybczynski theorem with a tech twist: Technical shocks impact output as follows:

1) Sign effects: Any tech shock that saves more A- than B-labour (i.e. $S_{QA} > S_{QB}$), creates more new firms/products in the A-intense sector than in the B-intense sector (i.e. $\hat{n}_1 > \hat{n}_2$). The opposite holds for $S_{QA} < S_{QB}$.
   a. This means that only the skill-bias of the tech shock matters for the sign of the relative output effect; the tech shock's sectoral composition is irrelevant.
2) Magnification effects:
   a. For **pure skill-biased shocks** (e.g. $\hat{a}_{j1} = \hat{a}_{j2}, j = A, B$), the standard Rybczynski effect applies without modification. The 'own sector' n_i expands more than proportionally, where 'own sector' means the sector intensive in the use of the factor experiencing the productivity gain. The other sector contracts.
   b. For **pure sector shocks** (e.g. $\hat{a}_{Ai} = \hat{a}_{Bi}$, i=1,2), the affected sector's output expands in the same proportion as the labour-savings shock, but the other sector's output is unchanged (e.g. $\hat{a}_{A1} = \hat{a}_{B1} < 0$ yields $\hat{X}_1 = -\hat{a}_{j1}, \hat{X}_2 = 0$).
      i. This result is obvious since a sector-specific tech shock expands the endowment in exactly the same proportions as the affected sector uses factors, so all the saved factors are employed in the progressing sector.
   c. For **Hicks Neutral shocks** (i.e. $\hat{a}_{ji} = \hat{a}_{ji}, j = A, B, j = 1{,}2$), all sectors expand in proportion to the shock.

Sector- and factor-specific shocks are a mixture of 2a) and 2b). See Appendix for the general analysis.

One sector- and factor-specific shock that is particularly relevant to the discussion is a skill-specific and sector-specific advance that boosts A-worker productivity in the A-intense sector 1. Plainly, $S_{QA} > S_{QB} = 0$, so the result is an output expansion of sector 1 than is more than proportionate to the tech shock, i.e. $\hat{n}_1 = \hat{X}_1 > \hat{a}_{A1} > 0 > \hat{X}_2 = \hat{n}_2$. **Result 1** tells us what happens to output, but what about jobs?

As it turns out, the magnification effect is critical. If jobs are to be created in a sector, the number of new firms has to expand more than proportionally compared to the labour-saving shock. To wit:

[2] This strict correspondence between technical change and endowment changes was introduced formally in Jones (1965 Section IX).

Result 2a: Rybczynski theorem with a jobs twist: The jobs impact of tech-shocks are marked by:

1) For **pure skill-biased shocks,** technical progress that raises productivity only of, say, A-workers creates new jobs in the sector that uses A-workers most intensely (sector 1); it destroys jobs in the other sector. NB:
    a. These new jobs involve producing new, previously unknown products, so we cannot know the names of the new jobs but we can know the sector in which they will appear.
    b. The job creation extends to both types of workers in sector 1.
    c. The shock destroys jobs for both types of workers in sector 2 since it leads to exit of sector-2 firms and elimination of their products from the market.

    Progress biased towards B-workers has the opposite effect.
2) For **pure sector shocks** technical change that raises productivity A- and B-workers equally but only in one sector produces no job creation or destruction in either sector.
3) For **Hicks Neutral shocks**, technical change that raises productivity A- and B-workers equally and in both sectors produce no job creation or destruction in either sector.

See the appendix for details.

A similar exercise using the Stolper-Samuelson theorem shows that real wage effects depend upon the sector in which the technical progress occurs, not the factor. This is exactly the opposite of the jobs impact where it was the factor that mattered, not the sector. In particular, pure sector-biased productivity advances raise the real wage of workers used intensively in the advancing sector; the other type of workers lose. A productivity gain that saves the same amount in both sectors, in the sense that $S_{C1} = S_{C2}$, will raise real wages in both sectors.

Result 2b: Stolper-Samuelson theorem with a tech twist:

1) For **pure skill-biased shocks,** technical progress that raises productivity only of, say, only A-workers is non-negative if $S_{C1} \geq S_{C2}$, and negative otherwise. Given the imposed symmetry, this holds, so A-workers gain. B-workers lose as per the usual Stolper-Samuelson logic. Progress biased towards B-workers has the opposite effect.
2) For **pure sector shocks** technical change that raises productivity A- and B-workers equally but only in one sector produces a real wage rise for the type of labour used intensively in the advancing sector; the real wage of the other type of labour falls.
3) For **Hicks Neutral shocks**, technical change that raises productivity A- and B-workers equally and in both sectors produce raises all real wages proportionally.

See the appendix for details.

With these results in hand, we turn to the discussion of how biased technological progress could account for the first grand structure transformation and the inequality results that followed it.

2.3 An interpretation of the results

If we think of manufacturing as sector 1 and A-workers as manual workers, the great technological innovations of the 18th and early 19th centuries were mechanisation of manufacturing. In the model that corresponds to very lopsided reduction in the amount of labour needed to produce any manufactured good, namely a_{A1}. To be concrete, and to clarify the economic logic, assume that the only thing mechanisation did was reduce a_{A1}.

What Result 2 tells us is that the shock will create jobs for A-workers in manufacturing and destroy jobs for manual workers in the other industry. The result is that the share of workers in manufacturing would rise. For a partial equilibrium perspective, it might seem odd that workers were moving into the sector where automation was advancing fastest, but GE considerations – totally apart from the classic price-lowering-demand-expanding effects – link automation with job creation. As shall become clear when we contrast this with a different type of technology in the next section, this result is entwined with the nature of technology shock.

2.3.1 Technology and the nature of work

During the shift of workers into factories in the 1870-1970 period, people had little idea of the names of the new jobs that would appear in the manufacturing sector. After all, all sort of products involving electricity, lighting, telecommunications, internal combustion engines, and industrial chemistry where invented well into the structural transformation. The light bulb, the radio, and the electric fan, just to take a few examples, were invented in invented in 1879, 1899, and 1882 (respectively). This is exactly how it happens in the model. The expansion of sector 1 involves the introduction of previously unknown varieties.

But while the names of the new jobs were unknowable, GE considerations shed light on the nature of the new jobs. Since the technology was sparing of manual labour, the new jobs had to be manual-labour intensive. Moreover, they had to be in the sector that used manual labour intensively.

3 Technology as better a substitute for manual workers and better tools for brain workers

The technology in the previous section was conceptualised in a standard way – sometimes called the 'factor augmenting' technical progress. This approach makes strong (implicit) assumptions about the role of humans in production. By channelling all technology improvements through human capacities, this approach makes it impossible, or at least very awkward, to take the human out of the production process. That is exactly why we got the somewhat counterintuitive result that the technology was inevitably job-creating in the sector where it advanced fastest. But in the 21st century is this still the right way to conceptualise technology?

There is an argument that the nature of technological advances changed qualitatively sometime around 1970 with the advent of computerisation – what some call the 'third industrial revolution'. Industrial applications of computers and telecommunications developed gradually from WWII, eventually reached into factories from the 1960s with numerical control systems. The year 1973 provides a convenient starting date for the 'services transformation' since that is the year that Texas Instrument employees Gary Boone and Michael Cochran patented the first 'computer on a chip'. This was revolutionary.

Putting a computer on a chip made earlier approaches to building computers obsolete; before, computers were built up from racks of circuit boards. By combining on a single thumb-nail sized device, the 'brain' (central processing unit, or CPU), digital memory, and circuits to handle inputs and outputs, the computer-on-a-chip reduced the cost and improved reliability – all while reducing power usage and thus overheating problems. And sticking a computer-on-a-chip into a robot arm allowed many repetitive mechanical tasks to be automated. Furthermore, the same robot could be quickly reprogrammed to do other tasks when the time came so the capital expenditure could be limited.

This was a watershed in the history of automation. It gave machines cognitive capacities that they never had. True, it was only the 'thinking slow' capacities – to use Kahneman's term for logical, conscious thinking – but it allowed a disassociation of machines and human that was previously unheard of; robots, for example, could weld together some cars parts without continuous human intervention. In short, the technology created better substitutes for humans who worked their hands.

Note that the automation focused on routine tasks (thinking slow) since these were the ones that humans were most successful at turning into computer code. Given the 20th century programming technology, 'think fast' tasks – like natural language processing, computer vision, computer reading, and the like – were beyond the scope of automation. Humans were still needed for these thinking-fast tasks.

What follows is a minimalist framework to clarify the economic mechanisms engaged by this very different form of technological progress. In short, it is an accounting of how computerisation led to a contraction of employment in the manufacturing sector and expansion of employment in the service sector. In the terminology of the seminal work by Acemoglu and Restrepo (2016), this section considers worker replacing technological progress.

3.1 A Ricardo-Roy model of robots and downgrade unemployment

The economy in this section closely resembles that of the previous section with a few complications and a few simplifications. There are two sectors, two types of skills, and two types of workers.

The first complication is allowing workers to have mixed skill sets. A-type workers are still relatively well-endowed with M-skills, and B-workers are still relatively well endowed with T-skills, but they each have some of both. A key assumption is that workers can only have one job and thus can sell the skills to only one firm.

To simplify the expressions, A-workers have 2 units of M and 1 unit of T while B-workers have the reverse (2 units of T and 1 unit of M), and factor intensity differences are taken to be extreme and polar opposites, namely sector-2 varieties are produced with only T-skills (one unit of T per sector-2 firm), while sector-1 varieties are made with only M-skills (one unit of M per sector-1 firm). Given the perfect identification of workers and skills in the previous model, this is equivalent to assuming $a_{A1} = a_{B2} = 1, a_{A2} = a_{B1} = 0$, but note that the worker-skill equivalence is broken here, so it is best to think in skill-space rather than worker-space.

Monopolistic competition now links prices to skill-rewards, which we denote with $v_i, i = M, T$, i.e.:

$$p_1 = v_M = 1, \qquad p_2 = v_T, \qquad \bar{x}_i = F_i\sigma; \;\; i = 1,2 \qquad (i)$$

The model is first solved for static technology, lining it up so most variables are unity initially.

3.1.1 Pre-robot equilibrium

Analysis begins with the worker's employment choices. Workers go to the sector that affords them the highest wage. If they work in the M-intensive sector, their T-skills will be useless; if they work in the T-intensive sector, their M-skills will be useless. Plainly, B-types get jobs in sector 2 if and only if $2v_T > v_M$. A-types work in sector 1 if and only if $2v_M > v_T$. Assortative matching (A and B types work in sector 1 and 2, respectively) occurs when:

$$2v_M > v_T > v_M/2 \qquad (11)$$

Once we know perfect sorting occurs, the rest of the equilibrium is straightforward. Workers' sectoral choice determines the inputs and thus the outputs of sectors 1 and 2. These, in turn, determine the relative price of goods and thus the v_i's. The equilibrium is closed by noting the rewards to skills determine workers' choice of sectors. Specifically, with sorting:

$$X_1 = M = 2L, \qquad X_2 = T = 2L \tag{ii}$$

since $L_A = L_B = L$. The numeraire choice, and the market-clearing and pricing conditions imply:

$$v_T = p_2 = \frac{X_1}{X_2} = \frac{M}{T}, \qquad p_1 = v_M = 1, \qquad w_A = 2v_M,\ w_B = 2v_T \tag{12}$$

where $M,\ T$ denote the supply of skills actually employed in equilibrium. The expressions for wages follow from workers' skill sets and sector choices. Using the symmetry:

$$p_2 = p_1 = v_M = v_T = \frac{w_B}{2} = \frac{w_A}{2} = \frac{\omega_A}{2} = \frac{w_B}{2} = 1$$

Note that a third of the economy's supply if each skill is unemployed in equilibrium; this 'skills unemployment' is essential to the analysis.

3.2 Computerisation: Industrial robots performing manual tasks

Technology is conceptualised as arriving in the form of industrial robots. For simplicity, robots appear exogenously, costlessly, and to be concrete, they are owned by B-type workers (who are now both workers and robot owners). Robots can supply only M-skills, so this is an extreme form skill-biased and sector-biased technological progress. Robot-generated M-skills are perfect substitutes for the M-skills of humans. The number of robots is fixed at N_R but we consider the impact of the robots becoming more proficient in the sense of providing more M-skills per robot, specifically:

$$M_R = A_R N_R, \qquad M = 2L_A + M_R \tag{\#}$$

where $A_R > 0$ is a factor-augmenting type of technology for robots ($dA_R > 0$ means better robots). The second expression holds in the sorting equilibrium. Note that we remain in the two-skill, two-sector setting even though there are three factors.

3.2.1 A three-phase impact

The fact that the sorting condition, i.e., $2v_M > v_T > v_M/2$, involves strict inequalities alerts us that the evolution of the economy will be marked by phases. In phase I, $2v_M > v_T > v_M/2$, continues to hold even as the rising supply of M-skills depresses its relative reward, thus making A-types closer to indifferent to working in the two sectors. At some point, v_T will have risen sufficiently to push $2v_M = v_T > v_M/2$. This happens when the robot supply of M-skills reaches M_R^i where $v_T = 2$, i.e.:

$$2 = \frac{2L_A + M_R^i}{2L_B}$$

so $M_R^i = 2L$. When M_R reaches this point, phase II starts.

3.2.2 Blue collars turn white

Phase II is marked by A-workers shifting into sector 2. The rate of the shift is such that the output of both sectors rises in tandem so there is no change in prices or rewards to skills during phase II. Direct calculation shows that the number of switchers (i.e. of A-type workers who have switched into their non-comparative advantage sector), denoted as L_{A2}, is:[3]

$$L_{A2} = \frac{M_R - 2L}{4}, \qquad R_M > 2L$$

Once all A-workers have changed sectors, phase III begins, and v_T continues its rise with $2v_M < v_T$, so all workers strictly prefer working in sector 2. This happens when M_R is sufficiently high to ensure $v_T = 2$ without any humans in sector 1. Labelling this threshold as M_R^{ii}:

$$2 = \frac{R_M^{ii}}{2L_B + L_A}$$

so $R_M^{ii} = 6L$. During phase III, the relative price of sector 2 products rises along with the real wages for both types of workers. To summarise:

Result 3 (jobs effects): *Technological progress which is sector- and factor- specific destroys jobs for workers in the sector experiencing the progress. New jobs involved in the production of new products are created in the other sector.*

Notice the stark difference from the impact of factor-augmenting technology in the previous section. Robots are a form of sector- and factor-specific technological advance but here the progress drives human out of the sector. Under the factor-augmenting technological conceptualisation, such progress pulled humans into the sector. Figure 1 illustrates the shift of workers and relative reward to the robot-competing skill. The level of M-skills provided by robots, M_R, is a proxy for the size of the technology shock.

[3] The expression is the L_{A2} that keeps $v_T = 2$, i.e. $2 = \frac{2(L-L_{A2})+M_R}{2L+L_{A2}}$ for $M_R > M_R^i, L_{A2} > 0$.

Figure 1: Three-phase structural adjustment to robotic technical progress

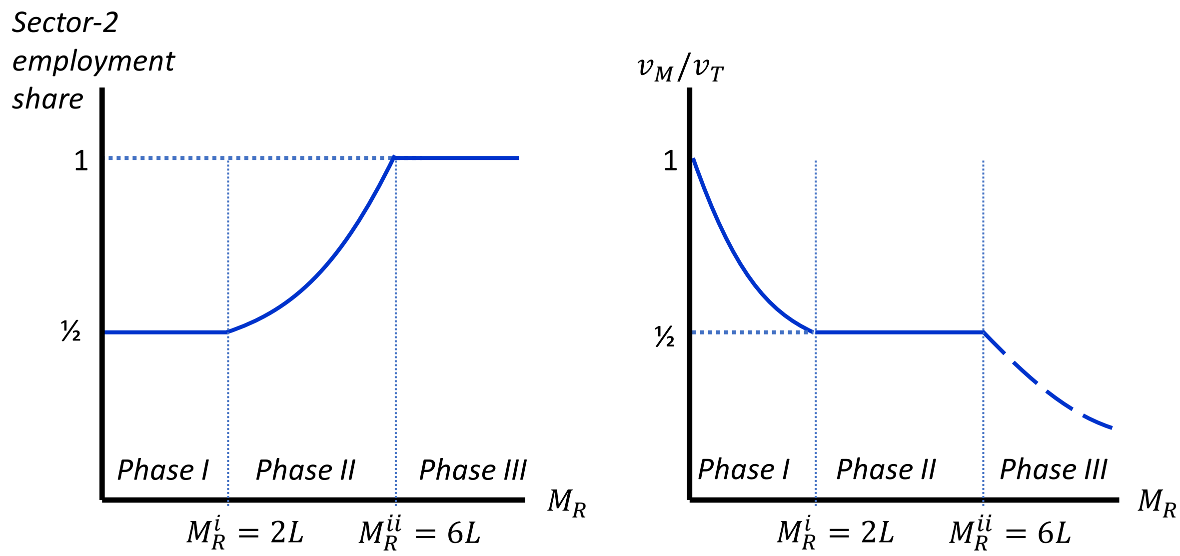

3.2.3 Real wages, incomes, inequality, and 'downgrade unemployment'

In all phases, the reward to M-skills is constant in terms of the numeraire (v_M=1), but in phase III only robots earn v_M. The reward to T-skills, v_T, rises as the price of sector-2 products rises in phase I and phase III. In real terms, A-workers start with real wages, $\omega_A = 2$, and proceed to see their real earnings fall as the relative price of sector-2 products rises (the price index equals $(v_T)^{1/2}$). B-workers start with $\omega_B = 2$, and see their real wages rise in phases I and III since $w_B = 2v_T$, rises faster than goods prices. Specifically:[4]

$$\omega_{A,I\&II} = \frac{2}{\sqrt{v_T}}, \quad \omega_{A,III} = \sqrt{v_T}, \qquad \omega_B = 2\sqrt{v_T}$$

using subscripts to indicate phases in an obvious notation. Observe that at both $R_M^i = 2L$ and $R_M^{ii} = 6L$, $\omega_A = \sqrt{2}$, and $\omega_B = 2\sqrt{2}$. These findings are reflected in Figure 2 and Result 4:

Result 4 (welfare effects): *The workers who are well-endowed with the skills that robots are providing (A-workers) are at first harmed by the technology while other workers are helped. However, once A-workers have moved fully out of the sector and are thus no longer directly competing with robots, further advances in robot proficiency are beneficial to all types of workers. Eventually, A-types see their real incomes rise above the pre-robot level.*

[4] The expressions follow from $\omega_{A,I\&II} = \frac{2v_M}{\sqrt{v_T}}$, $\omega_{A,III} = \frac{v_T}{\sqrt{v_T}}$, $\omega_B = \frac{2v_T}{\sqrt{v_T}}$,

Figure 2: Real wage evolution over the three phases

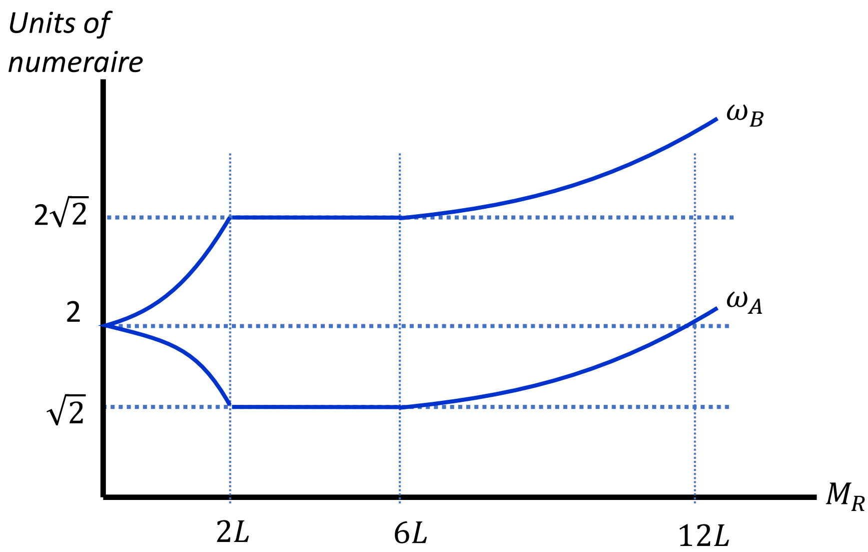

The value added that accrues to robot owners is simple to calculate. Robots provide only M-skills and these are, at any point in time, both valuable and limited, so perfectly competitive firms pay robot owners the same per-unit price of M-skills as they would pay to a human, namely $v_M = 1$. As the quantity of M-skills provided per robot rises steadily, the own-product wage rises, but the real value falls. As a consequence, in real terms, robot owners gain in line with A-workers.

In terms of GDP shares, the robots' share rises from zero to ½ during phases 1 and 2, but then stays at half due to the Cobb-Douglas preferences. Obviously, this means labour's share of GDP drops from 100% to 50% over the same timeframe. Wage and income inequality also rise.

With two income groups, the Gini coefficient equals the high-income group's share of GDP minus its share of population: $G = s_Y - s_L$, where s_Y and s_L are the B-workers' share of GDP and population respectively. Having allocated robot ownership to B-workers arbitrarily, B-workers' GDP share depends upon real wages and the reward paid to robot-owner. Due to the Cobb-Douglas preferences, half of national expenditure is devoted to sector-1 varieties and half to sector-2 varieties. With zero pure profits, this expenditure translates into factor payments, so the factors active in each sector get half the GDP in all phases. At the start of phase I, the Gini is zero since robots are getting nothing and wages are equal, but as robots start producing and earning more income inequality rises steadily due to both wage inequality (as discussed above) and due to the uneven distribution of robot income.

From the beginning of phase III onward, robots get half of GDP, and workers divide the rest with B-workers getting twice as much as A-workers. Thus the GDP share going to B-workers, s_B, is 5/6, leaving s_A equal to 1/6, while the shares are population are both ½. The Gini at the start of phases I and III are thus:

$$Gini_I = 0, \qquad Gini_{III} = 1/3$$

The Gini measure of income inequality rises steadily from the beginning of phase I till the start of phase III where it plateaus.

Redistribution of robot income could fully redress the rising inequality, since robots produce half of GDP. Consider the calculation for phase III where the Gini has stabilised at 1/3rd. Since B-workers get twice as much wage income as A-workers from sector 2, income equality could be restored by allocating 2/3rds of robot income to A-workers. This, in short, is an economy where robot incomes would be a means of avoiding the inequality aspects of technological progress. To summarise:

Result 5 (Rising inequality and the nature of future work):

1) *Income inequality as measured by the Gini coefficient rises as robots get more productive;*
2) *Redistribution of robot income could avert any rise in inequality.*

A-workers change sectors optimally, but by phase III none of them are using the skills in which they have an intrinsic 'comparative advantage'. For the A-type workers, the shift is essentially a downgrading of jobs – at least in terms of utilising their talents. They do not become unemployed, but they are unable to take advantage of their M-skills in their new jobs.

Note also that by phase III, robots have taken over all M-tasks in the economy – thus fundamentally transforming the nature of work.

Result 6 (nature of work and downgrading):

At the economy wide level, advancing robot technology changes the nature of work. From phase three onward, robots do what they can; humans do what robots cannot do. Better robots create no unemployment in terms of the number of workers with jobs, but it leads to the unemployment of the M-skills of all workers. This is not a new development for B-workers, but for A-workers this could be viewed as 'downgrade unemployment'.

Proposition 7 (labour share of GDP): *Throughout phase one, labour share of GDP falls from 100% to 50%, and stays there from the beginning of phase III onwards.*

Stepping outside the model, Phase II can be thought of as the de-industrialisation of advanced economies since the 1970s, namely the shift of jobs from factories to offices.

3.3 An interpretation of the shift of jobs to offices

If we think of sector 1 as manufacturing and A-workers as blue-collar workers, and sector 2 as the service sector and B-workers as white collar workers, this stylised account of the impact of robots provides a suggestive accounting of why technology had such different effects in the century before 1970 and the half century after 1970. Labour saving technology in manufacturing created jobs when it was labour-augmented but destroyed jobs when it was labour substituting.

3.3.1 Technology and the nature of work redux

When computerisation arose in the 18th century and accelerated at the end of the 19th century, people knew that technology would destroy many jobs. New jobs would, and were created, but at the time, people could not know the names of the jobs, or what sort of things workers would be making.

Some early visionaries talked about the post-industrial society, especially Bell (1974). He pointed out that by 1973, half the US workforce was in services, and he predicted that this expansion would continue and spread abroad since computerization was enhancing the productivity of service workers. Again, something like GE reasoning was being used to predict the nature of future work even if Bell had no idea of the particulars of the services that future workers would be producing. But the nature of the GE insights is quite different.

The GE implications for the nature of work in the previous section – where machine-helping-humans technology made farmers into factory workers – was based on the constraint that all the workers got reemployed. Rybczynski logic guided our thinking as to the factor-intensity of the new jobs and sectors in which they would appear.

In this sector – where robots-replacing-humans technology made factory workers into service workers – the GE hints for the nature of work are radically different. Here the nature of future work is predicted by a process of elimination; humans do what robots can't do. In this simple model, all manual tasks are eventually taken over by robots; the only tasks left for humans to do was thinking tasks. The GE insight then is that the new jobs would be intensive in the task that robots couldn't do.

In short, GE insights applicable to the 'blue collars turn white' transformation and robot technology are radically different than they were in the 'farm to factory' transformation driven by human-augmenting technology.

The nature of robots, however, is changing rapidly. They are starting to acquire the ability to provide thinking skills. This will surely displace many human jobs. The next section looks at what GE insights can tell us about where the new jobs will appear and what they will be like.

4 AI as a better substitute for brain workers and better tools for manual workers

Just as computerisation can be thought of as a technological watershed in the 1970s, Machine Learning can be viewed as a watershed from about 2016. A convenient signpost for this second technological watershed is the reversal of Moravec's Paradox.

In the late stone ages of AI, 1988 to be specific, AI-pioneer Hans Moravec wrote, "it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility." [5] Until recently, computers were good at the hard stuff but bad at the easy stuff. The deep reason was programming.

Humans taught computers to do things with computer programs. These explained, step-by-step, what the computer should do. This meant that before we could teach computers to think, we had to understand how we think, step-by-step. The paradox arose since, as another early hero of AI, Marvin Minsky put it, "we're least aware of what our minds do best." We understand how we do arithmetic, algebra, and archery; we haven't a clue as to how we recognize a cat, or keep our balance when running over hill and dale. Machine learning solved the paradox by skipping the programming.

With machine learning, the computer (the 'machine' part) estimates a very large statistical model of how to guess a solution to a particular problem (the 'learning' part). Thanks to advances in computing power, and access to big data, computers trained by machine learning routinely achieve human-level performance on specific thinking tasks - like recognizing faces or translating Chinese into English. In

[5] Moravec, Hans (1988), Mind Children, Harvard University Press.

essence, such software are 'white-collar robots', where the white collar part refers to the workers replaced by the tech.

This section explores the impact of technology when technology is providing better substitutes for thinking skills.

4.1 Machine Learning: Algorithms performing thinking tasks

The economy considered here is exactly as in the previous section, and the analysis starts from the phase III situation (where all humans are working in sector 2). The single addition is related to the development of commercially viable AI. Computer models, trained with AI algorithms (such as 'deep learning', neural networks, or similar), are able to provide thinking skills that are perfect substitutes with human-provided thinking skills. Thus, the total supply of T-skills in the economy is now:

$$T_R = A_T N_A, \qquad T = 3L + R_M \tag{\#}$$

where T_R is the robot-supplied T-skills, and, with a slight abuse of the technical terms, the robot providers of T-skills are labelled 'algorithms' with the number of algorithms denoted as N_A. The supply of AI-generated thinking skills is N_A times the effectiveness of each algorithm, A_T. As before, the technological progress is costless and exogenous, and these 'white collar robots' are owned by B-workers.

4.1.1 White collars turn blue

The introduction of AI shifts the economy into a new phase, phase IV. Assuming algorithms get better faster than industrial robots, T rises relative to M, and this drives up sector-2 production relative to sector-1. The attendant relative output shift lowers p_2 and thus v_T. At first (i.e. during phase IV), the progress of AI has no impact on the allocation of jobs since $v_T > 2$ and all workers earn more in sector 2 (recall the analysis opens during phase III from the previous model).

At some point, v_T is driven back down to a level (namely $v_T = 1$) where A-workers start switching back to sector-1 jobs. The threshold for this shift between phase IV and V is defined by:

$$1 = \frac{M_R^{iv}}{3L + T_R^{iv}}$$

so phase V starts when the quantity of AI-generated T-skills reaches $T_R^{iv} = M_R^{iv} - 3L$. As happened in phase II, the flow of A-workers from sector-2 to sector-1 proceeds at a pace that keeps $v_T = 1$. When all A-workers (who have 2 units of M-skills to sell) are in sector-1, the economy is back at a situation of assortive matching and a new phase, phase VI starts. Here $v_T = p_2$ is:

$$v_T = \frac{2L + M_R}{2L + T_R}$$

Thus the threshold between phases V and phase VI, is defined by

$$T_R^v = M_R^v$$

Assuming AI productivity continues to advance faster than industry robot productivity, v_T will continue to decline. Once $v_T = \frac{1}{2}$, B-workers, who have only 1 unit of M-skills to sell, start to switch to sector-1 as well. This is the start of phase VII, the threshold for which is defined by

$$\frac{1}{2} = \frac{2L + M_R^v}{2L + T_R^v}$$

so phase VII begins once $T_R^v = 2(L + M_R^v)$. Again, the reallocation of jobs flattens AI's impact on v_T, so $v_T = \frac{1}{2}$ while B-workers are transitioning to sector 2. The final phase, phase VIII, begins once all B-workers have made the transition. From this point, the price of sector-2 varieties continues to fall, bringing v_M down with it, but $v_M = 1$, as always (by choice of numeraire). Phase VIII starts when T_R surpasses:

$$1 = \frac{3L + M_R^{vi}}{T_R^{vi}}$$

so phase VIII is triggered when T_R exceeds $T_R^{vi} = 3L + M_R^{vi}$. The phases are illustrated in Figure 3.

Figure 3: Phases of structural adjustment to AI

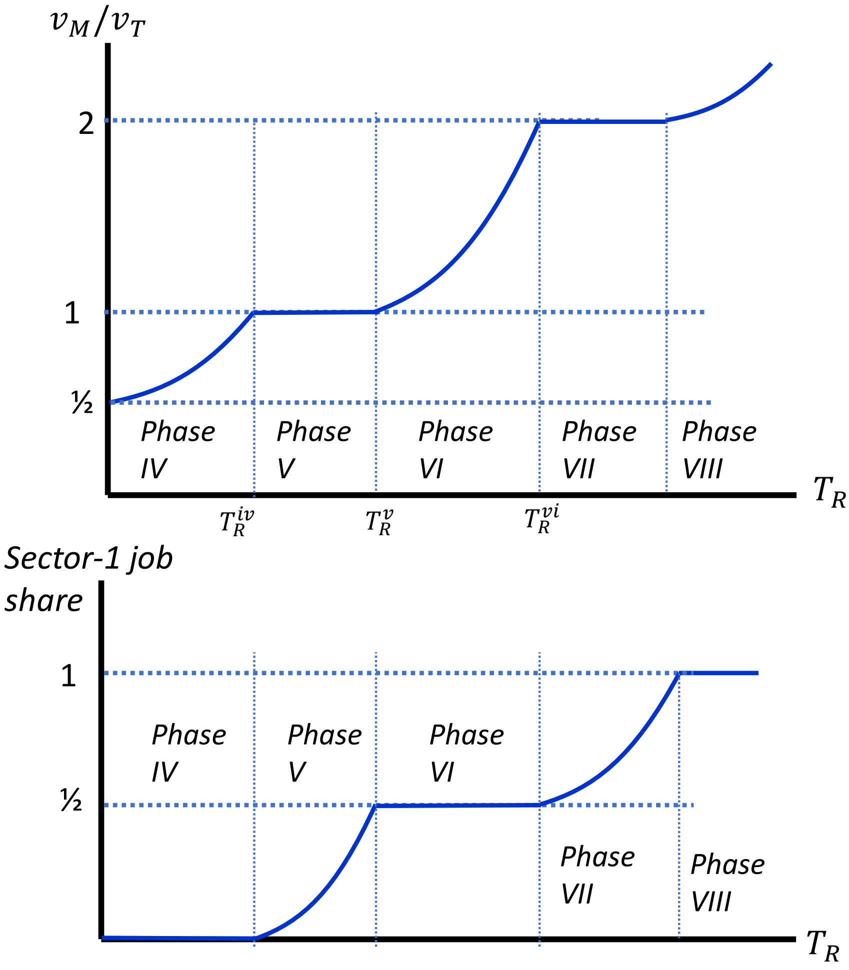

In terms of real wages, A-workers and B-workers are harmed by the drop in v_T during phase IV since their wage falls faster than prices. The owners of AI, by contrast, win during this phase since they earn v_T for each unit of T-skills produced by their algorithms.

In phase V, when A-workers are moving to sector 1, the wages are flat along with prices so only AI-owners gain. At the end of this phase, each worker is in his or her comparative advantage sector and the equilibrium shifts to phase VI. Here A-workers benefit since their wages are pinned at $2v_M = 2$, but the price index continues to fall. B-workers continue to lose at the same pace as in phase IV. From the start of phase VII, B-workers start to join A-workers in sector 1 even though they can only sell one unit of M-skills rather than 2 units of T-skills. The shifting of workers freezes prices and wages, so again only AI-owners gain. A- and B-workers neither gain nor lose. Once the shift to sector 1 is complete, phase VIII begins and both types of workers gain.

4.1.2 Inequality and labour share of GDP

The rising wage inequality that was associated with the rise of industrial robots is reversed by the rise of AI-trained algorithms in phase V. That is, both A- and B-wages equal unity in terms of the numeraire in phase V. Beyond that, AI drives inequality back up but this time in favour of the previously disfavoured A-workers since now B-workers are experiencing the downgrade unemployment. By phase VIII, A-workers earn twice as much as B-workers since they have twice as much M to sell.

5 Concluding remarks

There is an argument that the economy is on the cusp of a third grand transformation. The first was all about goods. It shifted people from one type of manual work (agriculture) to another (manufacturing). The second destroyed jobs for those who worked mostly with their hands (in factories) while creating jobs for those who mostly worked with their heads (in offices). Today's technological impulse is doing a bit of this, but with so few workers left in factories, the extra capabilities that industrial robots are getting from machine learning constitute more of a continuation than a transformation. The real transformative impact will be felt in the service sector. But what will the future jobs be like?

The first part of any likely answer is an admission of ignorance and caution. Automation displaced jobs in the nineteenth and twentieth centuries. Human creativity invented "needs" that we did not even know we needed. That's why many of us today work in jobs that would sound very strange to Charles Dickens in nineteenth-century London. The jobs were created in service sectors since they were the sectors that were shielded from automation. The same will almost surely happen again today. Jobs will appear in sheltered sectors. But what sort of jobs will these be?

We cannot know what new jobs will be, but by studying the competitive advantage of AI, we can say quite a bit about what sheltered jobs will look like in the future – a point I discuss at length in my book (Baldwin 2019). By studying the things that AI-trained robots can already do well, we can predict that the jobs that survive competition from AI and the new jobs that will be created are those that stress humanity's great advantages. Algorithms have not been very successful at acquiring social intelligence, emotional intelligence, creativity, innovativeness, or the ability to deal with unknown situations. Experts estimate that it will take something like fifty years for AI to attain top-level human performance in social skills that are useful in the workplace, like social and emotional reasoning, coordination with many people, acting in emotionally appropriate ways, and social and emotional sensing.

This suggests that the most human skills will be sheltered from AI competition for many years. The implication is as simple as it is profound. Humanity will be important in most of the jobs of the future.

That is the good news. The bad news is that millions, perhaps hundreds of millions of white-collar and professional workers will have to change jobs and perhaps retrain.

If the correct figure turns out to be in the millions and it is spread over a couple of decades, the transition is unlikely to trigger a politically significant backlash. But it the correct figure is in the hundreds of millions and is spread over a single decade, the economic and social upheaval could lead to a backlash of the type we saw in the 1800s and 1900s.

This may sound alarmist, but to put in context, I close with an extended quote taken from a speech given by the US labour minister (Secretary of Labor) James J Davis in 1927 – just few years for the greatest upheaval of all – the Great Depression – and the greatest backlashes of all – the rise of communism, fascism, and New Deal capitalism.

Excerpt from a 1927 speech by US Sectary of Labor, James Davis:

> *"Every day sees the perfection of some new mechanical miracle that enables one man to do better and more quickly what many men used to do. In the past six years especially, our progress in the lavish use of power and then harnessing that power to high speed productivity machinery has been tremendous. Nothing like it has ever been seen on earth.*
>
> *We must, I think, soon begin to think a little less of our wonderful machines and a little more of our wonderful American workers, the alternative being that we may have discontent on our hands.*
>
> *Understand me, I am not an alarmist. If you take the long view there's nothing in sight that gives us great concern. I am no more concerned over the men once needed to blow bottles then I am over the seamstress that we once were afraid would starve when the sewing machine came in. In the end, every device that lightens human toil and increases production is a boon to humanity. It is only the period of adjustment, when machines turn workers out of their old jobs and into new ones, that we must learn to handle them so as to reduce distress to the minimum.*
>
> *We must ever go on, fearlessly scrapping all methods and old machines as fast as we find them obsolete. But we cannot afford the human and business waste of scrapping men. In former times the man suddenly displaced by a machine was left to his fate.*
>
> *The new invention we need is a way of caring for this fellow made temporarily jobless. In this enlightened day we want him to go on earning, buying, consuming, adding his bit to the national wealth in the form of product and wages."*

If only President Coolidge had listened to his minister!

References (incomplete)

Acemoglu, Daron & Autor, David. (2010). Skills, Tasks and Technologies: Implications for Employment and Earnings. National Bureau of Economic Research, Inc, NBER Working Papers. 4. 10.1016/S0169-7218(11)02410-5.

Bell, Daniel. The Coming of Post-Industrial Society. New York: Harper Colophon Books, 1974.

Hémous, David and Morten Olsen (2014). “The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality”, CEPR DP 10244, https://cepr.org/active/publications/discussion_papers/dp.php?dpno=10244

Herrendorf, Berthold & Rogerson, Richard & Valentinyi, Ákos, 2014. "Growth and Structural Transformation," Handbook of Economic Growth, in: Philippe Aghion & Steven Durlauf (ed.), Handbook of Economic Growth, edition 1, volume 2, chapter 6, pages 855-941, Elsevier.

Matsuyama, Kiminori (2009). Structural Change in an Interdependent World: A Global View of Manufacturing Decline," Journal of the European Economic Association. 7 (April-May 2009): 478-486.

Ngai, Rachel and Christopher Pissarides (2007). "Structural Change in a Multisector Model of Growth," American Economic Review, vol. 97(1), pages 429-443, March.

Xu, B. (2001). ‘Factor bias, sector bias, and the effects of technical progress on relative factor prices’, Journal of International Economics, 54(1), 5-25

Xu, Bin (2001). “Factor Bias, Sector Bias, and the Effects of Technical Progress on Relative Factor Prices,” Journal of International Economics, Volume 54, Issue 1, June 2001, 5-25.