Technology and the Future of Work:

Comparing the Potential Impact of AI Industrialisation on Labour Markets in Liberal and Coordinated Market Economies

By

Nikita Sidorov

Submitted to

Central European University

Department of Political Science

In partial fulfilment of the requirements for the degree of Master of Arts in Political Science

Supervisor:

Dr Imre Gergely Szabó

Vienna, Austria

2023

Author's Declaration

I, the undersigned, Nikita Sidorov, candidate for the MA degree in Political Science declare herewith that the present thesis is exclusively my own work, based on my research and only such external information as properly credited in notes and bibliography. I declare that no unidentified and illegitimate use was made of the work of others, and no part of the thesis infringes on any person's or institution's copyright. I also declare that no part of the thesis has been submitted in this form to any other institution of higher education for an academic degree.

Vienna, 07 June 2023

N. Sidorov

Signature

Abstract

The increasing integration of Artificial Intelligence (AI) into various spheres of society is a topic of significant interest, given its pace of development and deployment. Therefore, this study analyses the potential impact of AI on labour markets in Liberal (LMEs) and Coordinated Market Economies (CMEs), focusing on the United States and Germany. It extends theoretical assumptions of the Varieties of Capitalism scholarship by postulating that due to different innovation patterns and skill systems, jobs in LMEs and CMEs would differ in their exposure to automation. Methodologically, it utilises quantitative text analysis creating indices of AI-exposure of jobs by comparing US and German job tasks against titles of AI-related patents. The findings confirm the hypothesis and reveal a significant difference in the mean automation index between the two countries, with higher exposure to AI in the US. Furthermore, jobs requiring high and/or specific skills, such as engineers or doctors, in both economies are among the most exposed to AI, reinforcing the dominant theories suggesting that white-collar workers are most exposed during the 'new wave of automation'. Finally, this research underscores the necessity for increased attention to AI-driven automation, considering the plethora of potential social, political, and economic changes it may bring about.

Acknowledgements

First and foremost, I owe my deepest gratitude to my research supervisor, Dr Imre Gergely Szabó, whose guidance, feedback, and suggestions have been invaluable throughout my research. Moreover, I would like to express my appreciation for Imre's courses "Welfare States in the Current Era: Origins, Issues and Challenges" and "The Future of Work: Challenges for Politics and Policy." They were among the most insightful and helpful courses that I took, both for my academic research and daily life.

I am also grateful to Dr András Gál with whom I had an immense pleasure working the last two years. I would not be able to make it to CEU without his mentorship and support. My special thanks are extended to Robin Bellers from the Centre for Academic Writing for his meticulous proofreading of this and other academic papers during my studies. Furthermore, I am deeply grateful to the CEU student centre, especially to Maria Andrea Urbas and Andrea Borodi. Their immense assistance and support concerning my visa were invaluable. I am also thankful to Monika Zay for her patience and guidance through the labyrinth of CEU administrative processes. Also, this endeavour would not have been possible without the generous financial support provided by CEU itself, facilitating my studies.

My friends, within and beyond CEU, have been an extraordinary pillar of support in both academic and personal matters. Special mention goes to Marina Romanova and Vladislav Shatilov for being there in the most needed times. I am also indebted to Elena Lesnikova, Anastasia Kuprishkina, Egor Balalykin, Vasileios Zivas, Artur Baranov, David Avagyan, and many others for the cherished memories we have created over this and previous years. I am particularly thankful to Alexander (Sandro) Ketsbaya for his immense support during my time in Tbilisi and for all the khachapuris we shared together. His friendship made my exile significantly more bearable.

Finally, I would like to express my deep appreciation to my parents, Svetlana, Igor, and Luidmila. Their unconditional and everlasting love, support, and encouragement have been the foundation upon which all of this was possible. Thank you. Needless to say, all the compelling aspects of this work are due to the effort and advice of the people above; however, any errors or omissions found within are solely my responsibility.

Table of Contents

List of Figures

Introduction

In recent years, there has been a rapid transformation in the mechanisms of production, consumption, and employment. In the past half-century, post-industrialisation has led to the decline in industrial low-skill, high-paying jobs in the Western world (Pierson 2002). The gig economy, which provides limited social protection and benefits, has seen an expansion in popularity and currently employs more than 30% of US workers (Gallup 2018). Moreover, global supply chains have expanded to cover the entire planet, creating a highly interconnected and complex global economy.

Yet, the world is again standing on the precipice of another possible revolution. The development and industrialisation (i.e., mass commercial use) of computer systems capable of performing tasks that traditionally required human intelligence, such as understanding natural language, recognising patterns, and solving problems, introduce a variety of new challenges for modern polities. They range from the need to tailor educational policies to meet the evolving demands of the labour market, to addressing ethical concerns over data privacy and potential biases in AI systems. However, the focus of this paper is on the *potential impact of AI on the working environment; specifically, on examining the exposure of different jobs (through their respective tasks) to AI automation* - that is, the potential replacement of human labour with AI-powered solutions.

It is presumed that not all job tasks are equally susceptible to AI automation. For example, while high-pay high-skill occupations (e.g., 'white collar' technical workers, statisticians, and engineers) are being most-exposed to AI industrialisation (Felten, Raj, and Seamans 2019; Webb 2019), jobs involving communication, social interaction, and human physique are among the least-exposed, notwithstanding their skill level, – given the difficulties of automation of communication and replication of human movements.

At the same time, *Varieties of Capitalism* (VoC) literature postulates that due to a number of factors, different types of political economies vary in their technological innovation patterns: Liberal Market Economies (LMEs), like the US or the UK, are *radical* in their innovation, creating products that disrupt the current production processes, while Coordinated Market Economies (CMEs), like Germany or Japan, are *incremental*, and innovate by improving on the current production lines (Hall and Soskice 2001). One of the key factors explaining this difference is a *skill system* – the institutional arrangements and practices surrounding the acquisition, utilisation, and deployment of skills on part of the labour force. It is postulated that LMEs are characterised by the prevalence of *general* skills – broad abilities that are transferable throughout different employers/industries. Contrary to this, in CMEs *specific* skills dominate – they are narrow and unique to the particular employer/industry (Estevez-Abe, Iversen, and Soskice 2001; Hall and Soskice 2001).

Hence, while the current scholarship has mostly focused on the prospect of job automation in the US (Frey and Osborne 2013; Webb 2019), between OECD countries (Arntz, Gregory, and Zierahn 2016), or on the increasing mismatch in skills that are being replaced on the labour market, and the ones employers are expecting workers to have (Zarifhonarvar 2023); it lacks the 'theory-testing' comparison of job automation prospects between countries with different innovation patterns, i.e., between LMEs and CMEs. This study addresses this gap in literature by asking the following research question (RQ): *considering their distinct innovation patterns and skill systems, does the proclivity for job task automation diverge between LMEs and CMEs?*

I hypothesise that due to the differences in innovation patterns and skill systems, more jobs can be automated in the US than in Germany. Indeed, it seems that AI can put further strain on jobs requiring specific skills, given their hypothetical susceptibility to automation. To answer this RQ, I relied on *quantitative text analysis*. In particular, I extracted the texts of job tasks for the US and Germany (following the original VoC comparison), and texts of AI-related patents from the United States Patents and Trademark Office (USPTO). Then, I constructed an automation index for different jobs in the US and Germany and compared them. The findings confirm the predictions established in the literature. The prospects for automation in the US and Germany are quite different with the US jobs, on average, having an AI-exposure index of 0.02% vs 0.0051% for German jobs. The difference is also statistically significant, though when working with small numbers statistical significance may be distorted. At the same time, the countries seem to converge on the 'types' of jobs that are most and least exposed to automation: high-, in particular, specific-skill and non-routine jobs, such as engineers or doctors are most exposed to automation, while low-skill, routine roles or those requiring physical exertion, such as fitness trainers, are least exposed. However, the high exposure of some professions, such as doctors, warrant a further examination of the viability of the chosen methodology in calculating the exposure index given the obvious presence of communicative aspect in such occupations and thus a lower expected exposure rate.

Nevertheless, the analysis is warranted to provide insights both to scholars (to have a better grasp on social and economic implications of continuing automation) and policy-makers (to tailor educational policies to the changing demands). Additionally, given the severity of the possible negative political consequences of automation, such as technological unemployment, resentment, and de-gentrification, it is crucial to pay additional attention to the question of job automation in time, so that appropriate strategies can be devised to mitigate potential social and political challenges.

The paper is structured as follows: firstly, I introduce the reader to the debates regarding 'the future of work' examining possible benefits and negative effects of continuing automation and evaluate the potential political implications of automation. Then, I discuss the main theoretical framework of the paper – Varieties of Capitalism. I outline the differences in the institutional arrangements between LMEs and CMEs focusing on their propensities to innovation. Secondly, I bring forth my theoretical contribution to the field, comparing the impact of previous waves of automation, while hypothesising about the implications of the ongoing AI wave, on diverse skill systems. The comparison and hypothesis aim to provide a deeper understanding of the evolving dynamic between technology and labour in the era of AI. In particular, I postulate that since the introduction of laptops and Internet, the importance of general skills has been increasing. Then, I present the reader with my research question and hypothesis, asking whether the prospects of automation differ between LMEs and CMEs given their divergent preference for skill systems. Fourthly, I overview the methodology and delve deeper into the technicalities of text analysis methods employed. Then, I present he findings for the US and Germany and discuss them. I conclude by reminding the reader of a broader societal impact of automation and the necessity to take appropriate steps in time.

Chapter 1. The Future of Work

Given the complexity of the subject at hand, I divide this chapter into two main parts. Firstly, I discuss the phenomenon of job automation, emphasising both the potential benefits and pitfalls, such as the possible liberation of labour and the introduction of more meaningful jobs vs. technological unemployment and further job polarisation. I also put back the political in 'political economy' by highlighting the potential effect of job automation on the institutions of liberal democracies. The second section establishes the main conceptual framework of this work derived from the Varieties of Capitalism (VoC) scholarship. I argue that while the empirical evidence regarding the differences in innovation patterns between LMEs and CMEs is ambiguous, it is nevertheless a useful tool in differentiating the institutional arrangements in these economies.

1.1 The Future of Work & AI Automation

The future of work has long been a subject of great interest, with a focus on the potential impact of automation driven by technological advancements (be it a conveyor belt or an AI algorithm) on skills in demand in labour markets, political systems, and human lives more generally. Thus, in 1930, Keynes famously predicted that by 2030 people would work only 15 hours per week (Pecchi and Piga 2008) given the massive increases in productivity that technology brought about.[1] For better or worse, his predictions have not materialised. However, the same question arises once again with the claims that AI could replace a handful of jobs and thus people would be finally free to pursue their dreams.

[1] Moreover, a substantial part of the Keynes' "Economic Possibilities for our Grandchildren" essay was dedicated to the problem of the free time that would suddenly appear and the question of how people would occupy themselves.

In this context, it becomes crucial to consider the perspectives of those advocating for automation. The proponents of automation argue that machines would complement and assist human workers, making them more efficient by increasing their productivity[2] – the so-called 'enabling effect' (Frey 2019). To illustrate, the advent of spreadsheets in the 1970s presents an example of an enabling technology, as it did not make any job redundant, but instead made analysis easier and removed some routine tasks (Pethokoukis 2017). Furthermore, the proponents assume that automation would only touch on some of the *tasks* that people perform on their jobs (Autor 2015). For instance, while laboratory testing and assessment of the results by a nurse can theoretically be automated, the inter-personal component of interactions with patients cannot, given how clumsy the machines are in communications and menial labour (Lane and Saint-Martin 2021). Hence, this complementarity would lead to a "bigger pie effect," (Skidelsky 2023) when increased productivity results in an enlarged supply of goods and services, thus rendering long working hours unnecessary. Additionally, technological advancements may drive the development of new goods and services, creating new demands and thus generating novel job opportunities while enhancing the quality of life for consumers (Mokyr, Vickers, and Ziebarth 2015, 45). Lastly, the combination of these factors would mean that people have a greater choice between work and leisure. In essence, '*labour*'[3] would be

[2] The increase in productivity, however, leads to higher opportunity costs of not working: if it currently takes T hours to complete X additional tasks, resulting in a Y increase in compensation (as a premium, for example); under a more productive system, one could accomplish 4X tasks in the same T hours, theoretically resulting in a 4Y reward. *Consequently, the higher pay rate serves as an opportunity cost of not working the same T hours*. Nonetheless, as I will demonstrate later, not all jobs can be automated to the same extent. Therefore, jobs with greater potential for automation will have higher opportunity costs, and people will continue to work the same hours. Empirical studies seem to support this hypothesis, showing that individuals with tertiary education experienced a significantly smaller increase in leisure time compared to those with less than a high school degree, who gained an almost tenfold increase in leisure time (Aguiar and Hurst 2007).

[3] The distinction between lousy and lovely jobs resembles the differences between 'work' and 'labour' envisioned by Hannah Arendt (Arendt and Canovan 2010). She connected 'labour' with tasks providing for survival and satisfaction of biological needs. In essence, 'labour' is a menial activity that people do to earn a living, akin to a "bullshit job" in the words of Graeber (2019). On the other hand, 'work' is meant to leave an impact – an artefact – in the world, thus having a creative component to it.

replaced, de-commodified (Esping-Andersen 1990, 37), and people would '*work*' for the sake of achieving excellence in the craft, helping the community, socialising, and/or attaining social recognition (Gheaus and Herzog 2016).

At the same time, critics of automation contend that the current wave of automation, characterised by advanced AI and robotics, is different from previous ones (this issue is explored more thoroughly in sections 2.1 – 2.3), as it now includes even the automation of *mental* work, leaving few job tasks that robots cannot perform (Brynjolfsson and McAfee 2014, chap. 2). This may result in technological unemployment whereby technological solutions replace human labour. This outcome is made more likely due to the aforementioned susceptibility of non-routine tasks: Felten, Raj, and Seamans (2019) designate statisticians, engineers, programmers, and other 'white-collar' technical jobs as susceptible to automation. Additionally, middle-class jobs with *routine* tasks are also threatened by automation, thus increasing the severity of potential technological unemployment. Overall, while there is no consensus on the net-effect of automation, the optimists argue that it may finally free labour or, at least, automate the menial labour while people would continue to work more meaningful jobs. Contrary to this assumption, pessimists stress that the new wave of automation is different and would likely lead to a form of technological unemployment which, without adequate regulatory mechanism, may precipitate negative social consequences of increased inequality and poverty. While it is hard to predict, which effects would be more pronounced in the future, it is instrumental to discuss not only the 'economic' effects of automation but also its political ramifications.

1.2 Political Consequences of Automation

The debates surrounding the future of work necessarily have a political component to them as people spend most of their time in jobs which seem to influence their political preferences (Kohn 2001; Kitschelt and Rehm 2014), hence it is vital to discuss the political implications of automation.

Modern Western liberal democracies were built around the middle-class that emerged during the 'Golden Age' of capitalism lasting from the end of the Second World War into the oil shocks of the 1970s. During this time, the growth of the middle class was driven by the expansion of manufacturing industries and mid-skilled clerical occupations in the growing service sector, which provided well-paid, low-/medium-skill stable jobs for a large proportion of the population (Pierson 2002). This produced an upward social mobility that had facilitated the integration of the working class into the mainstream political system and helped foster a consensus around the welfare state, economic growth, and social equality (Held 2006). During this time, the working-class has undergone a process of *embourgeoisement* (gentrification) through which it has converged with middle-class values, aspirations, lifestyles, and structures of social relationships (Goldthorpe and Lockwood 1963). However, the neoliberal turn had begun to erode social cohesion due to exacerbated income inequalities (Hacker and Pierson 2010) and the condition of "permanent austerity" (Pierson 2002). This, coupled with further automation which targets exactly the middle-class jobs may further undermine the social consensus and lead to the effective disenfranchisement of large segments of the population, who might perceive the political systems as corrupted by the elites or attribute their misgivings to globalisation or immigrants (Dehesa 2006). Indeed, the scholarship shows that the theoretical 'losers of automation' have a higher potential of voting for the radical right-wing parties (Im et al. 2019) that may want to undermine democratic institutions.

The issue of automation and its impact on the future of work is complex, with both positive and negative consequences of further automation. While it offers potential benefits such as increased productivity, new job opportunities, and greater choice between work and leisure, it also raises concerns about job displacement, income inequality, and precarious employment. The political consequences of automation, particularly the rise of resentment politics and possible *de-embourgeoisement*, should not be overlooked. As humanity continues to advance technologically, it will become crucial to address these challenges through thoughtful policy-making. Having examined the debates surrounding the future of work as a context of this study, I turn to the analysis of its main conceptual underpinning for the empirical analysis – Varieties of Capitalism.

1.3 Varieties of Capitalism, Innovations, and Institutional Advantage

As this work employs the conceptual framework developed by Hall and Soskice (2001)[4] it is imperative to examine it in greater details. In their seminal work, H&S coined the concept of Varieties of Capitalism (VoC). They postulated that capitalism is not uniform, and while this approach is not new per se (Shonfield (1969)), it broke ground in connecting microeconomic principles of firms' operations to macroeconomic and political institutional arrangements. Additionally, they introduced two 'ideal types' of capitalist systems: Liberal Market Economies (LMEs) and Coordinated Market Economies (CMEs). Essentially, they posited the way firms in these two regimes solve the problems of transaction costs and informational asymmetry influences and reinforces the institutional environment under which they operate. I employ VoC as a conceptual foundation for its clear-cut distinction between different types of

[4] Hereinafter abbreviated as **H&S**.

capitalism and the divergence of skill systems therein that allows to better measure (and compare) the impact of AI on jobs, by using skills as an intervening variable through which AI impacts labour markets. In this section, I briefly analyse the differences between the two models emphasising their divergent propensities for technological innovation.

Firstly, the *labour market structures* differ between the economies in terms of flexibility, wage coordination, employment protection, and skills formation processes. LMEs, such as the United States or the United Kingdom, tend to have more flexible labour markets, with lower levels of employment protection, decentralised wage bargaining, and a greater reliance on market mechanisms for determining wages and working conditions (Esping-Andersen 1990; H&S). These factors, in turn, influence the skills that employees decide to perfect. Thus, low levels of employment security coupled with the mediocre unemployment security, characteristic of LMEs, lead to the emphasis on general skills, i.e., those that are not connected with a particular firm or industry (like problem-solving and computer literacy), which facilitate easy labour transition and flexibility between firms and industries (Estevez-Abe, Iversen, and Soskice 2001). Such flexibility, apart from producing higher levels of inequality, leads to the distinct way of technological transfer that is accomplished through the 'poaching' of employees – the transfer of employees to jobs with more lucrative conditions (including from the public to private sector). In contrast, CMEs, like Germany and Japan, feature more regulated labour markets, with higher levels of employment protection, centralised or coordinated wage bargaining, and stronger ties between employers and employees.

Therefore, the difference in labour market structures often results in more stable employment relationships and lower income inequality in CMEs compared to LMEs (Kenworthy 2001); given the rigidity of labour markets coupled with a prevalence of long-term contracts, and effective vocational training, employees have the incentive to invest in firm- and industry-specific skills.[5] These factors, however, can also limit the flexibility and adaptability of enterprises during downturns. Nonetheless, they allow firms to preserve the workforce, as patient capital diminishes the likelihood of employees facing terminations at the onset of an economic downturn. Regarding the innovations in CMEs, as I have postulated earlier, they rely on the constellation of inter-company cooperation, common standards setting, and subsequent diffusion of technology across the economy (Soskice 1997).

Secondly, in terms of *financial systems* LMEs are characterised by a high degree of market relations between the actors. This, together with widespread mergers, acquisitions, and takeovers, fosters a focus on profitability and leads to a lack of *trust* and cooperation between the firms decreasing the potential for joint Research and Development (R&D) venues. These economies are also highly financialised, boasting well-developed stock markets that generate substantial amounts of easily available venture capital (VC) funds (Black and Gilson 1998), which are essential for the innovative start-up companies. On the other hand, CMEs have a comparatively trusting environment between the companies partially induced by the "institutional complementarity between the legal system and the system of business coordination" (Casper 2001, 415) discouraging hostile actions and thus cultivating a favourable environment for joint R&D projects (H&S). Consequently, the availability of high-risk, high-reward VC is often limited in CMEs due to the prevailing bonds between banks and firms, a general bias towards debt financing, and labour market rigidities (Jeng and Wells 2000). Instead, CMEs tend to rely on 'patient capital' which offers a long-term investment approach that aligns with the gradual growth patterns of firms. For instance, a network of publicly owned

[5] Note, however, that Estevez-Abe, Iversen, and Soskice (2001) themselves observe a further division among countries that have both high employment and unemployment protection (such as Germany) and those which only have the former (Japan, Italy). Employees in the latter economies tend to rely mostly on firm-specific skills.

Sparkassen or *Landesbanken* in Germany significantly profited from providing low-cost capital to the *Mittelstand*, the small and medium-sized enterprises (SMEs) that form the cornerstone of the German economy (Howarth and Quaglia 2014).

Thus, firms, through their operational strategies and practices, play a significant role in structuring the institutional environment of an economy. Indeed, many differences can be observed between LMEs and CMEs but what is crucial in this analysis is that not only the ways of innovation but also the kind of innovations promoted by these economies differ. Therefore, as H&S (39) postulate, LMEs are characterised by *radical* innovation – "substantial shift in product lines", while CMEs demonstrate *incremental* innovation patterns – "continuous but small-scale improvement to existing product lines" (H&S, 39). Essentially, they connect the type of innovations with what they deem a 'comparative institutional advantage' of such a system. Thus, high market competition, easy access to VC, and flexible market structures promote inventions in technologically intensive spheres, such as semiconductors industries, bio-medical technology, and Information and Communication technology. At the same time, CMEs, due to the opposite institutional environment, excel at engineering or transport technologies that are more well-established. H&S tested their assumption using the data from the European Patent Office and, indeed, found the differences between the US and Germany in the relative share of patents registered by the country in different technologies (H&S, 42).

Although this theory does account for the locations and products of clusters like Silicon Valley and Baden-Württemberg, it overlooks several crucial factors. Firstly, the simple count of patents that H&S used may hinder the real picture as it does not capture the *importance* of a patent. For example, before 1982 in the rapidly growing field of Computed Tomography "almost half the patents [were] never cited" (Trajtenberg 1990, 181) thus probably not representing a breakthrough/radical innovation that is attributed to them by H&S. Secondly, H&S did not provide any explanation other than institutional effects on innovation: for instance, the immense role that the state and military have played in the rise of Silicon Valley and universities, such as Massachusetts Institute of Technology or Stanford (Leslie 1993). Also, Taylor (2004) has found that the specialisation of LMEs in radical innovation depends on the inclusion of the US in the equation as it is an outlier, otherwise, the H&S findings were disconfirmed. Later, Akkermans, Castaldi, and Los (2009) established that patterns of innovation (radical vs incremental) are industry-dependent: LMEs radically innovate in the fields of petrochemical products and electronics, while CMEs are 'radical' in transport and machinery industries. However, when innovation is measured by the originality of the invention, i.e., by the number of technologies that facilitated the production of innovation, the H&S hypothesis actually holds with "inventors in LMEs draw[-ing] on a much broader base of technologies in producing new innovations" (Akkermans, Castaldi, and Los 2009, 189), thus partially confirming the original findings. Collectively, these factors enabled scholars to, if not outright disprove, at least weaken the original findings.

Overall, I have explored the differences between LMEs and CMEs in terms of their labour market structures (skills promoted) and financial systems that underline them. Furthermore, while criticism of VoC empirical findings holds merit, it is at least theoretically plausible that different types of economies would innovate differently and while the application of the theory requires precision given the plethora of nuances when dealing with the patents (discussed in the Methodology section), some of the differences between LMEs and CMEs can be observed in the field of innovation. Now, I delve deeper into the topic of skills and theorise about the different impact of the three waves of automation on general and specific skills.

Chapter 2. Theoretical Framework: Waves of Automation

Considering the pivotal nature of the general vs. specific skills divide for VoC theory it is important to distil the potential impact of AI not only on jobs but also on skill systems in order to assess the prospects of automation in LMEs and CMEs. Additionally, the comparison with previous waves of automation would help to contrast the effect of the current 'AI revolution' with the introduction of mechanised processes and standardised production (Industrial Revolution)[6] and early digitalisation through computers and the Internet (Digital Revolution). Consequently, this chapter examines the impact of the three waves of automation – Industrial Revolution, Digital Revolution, and AI revolution, on skills. After discussing the theoretical part, I introduce the reader to the main Research Question and hypothesis of this study.

2.1 Industrial Revolution

The first wave of automation, typified by mechanisation during the Industrial Revolution, exerted a transformative impact on the skill set demanded from the workforce. It was characterised by the *replacement of manual labour with machinery and the shift to mass production, affecting jobs requiring specific skills*. Artisans, who epitomised the specific skillset (e.g., woodwork, textile manufacturing) found their roles progressively marginalised with the rise of standardised and mass-produced goods. For instance, the textile industry, which was originally dominated by skilled weavers and spinners, underwent a radical transformation due to the invention of the power loom and the spinning jenny that mechanised weaving and spinning, thereby 'de-skilled' the artisanal jobs (Mokyr 1992).

[6] Note that while the mechanisation and standardisation of production are distinct processes, they are analysed together under the umbrella term 'Industrial Revolution' given their (relative) co-occurrence and the compound effect on skills demanded from labour.

Additionally, the transformation of the automobile industry serves as a salient example of the shift from craft production to mass production. As noted by Womack (2007), this transition was marked by a stark change in several key characteristics of the production process. Under the craft production regime, the workforce primarily consisted of highly skilled artisans who improved their skills through apprenticeship. Production was decentralised, with general-purpose equipment being used across diverse tasks. The output, while of high quality, was limited in volume. This paradigm underwent a radical shift during the change to mass production. Skilled artisans were replaced by low-skilled labourers, each assigned a specific task in an assembly line. Opportunities for skill advancement were minimal, as each worker's role was confined to a single, repetitive task. Furthermore, under then popular theory of scientific management of Taylorism, the focus was placed on task efficiency rather than skill development, as Taylor advocated for reducing complex work into simpler, more manageable tasks. This mechanistic approach left little room for personal growth or skill advancement, as workers became confined to executing specific, repetitive tasks (Braverman 1998). Production became centralised in factories equipped with highly specialised machinery, designed to be simple in operation. This shift allowed for the exploitation of economies of scale, leading to a substantial increase in output. As a result, goods (such as automobiles) could be produced more efficiently and cost-effectively, making them more accessible to the broader public.

To summarise, the Industrial Revolution, had a significant impact on the labour market, with several key outcomes: 1) artisans and craftsmen, who were the bearers of specific skills, found their roles greatly diminished. These individuals, who once dominated industries such as textiles and metalwork, were mostly relegated to niche markets, typically in the luxury segment (Womack 2007). Unable to compete with the lower prices and larger scale of mass production, many of these skilled workers faced economic marginalisation. 2) Conversely, the demand for unskilled labour surged. Factories required a vast workforce to operate machinery and maintain production lines. This included the employment of *strike-breakers*, who were brought in to replace permanent workers during strikes. However, this high demand for unskilled labour came with its own set of challenges. Turnover rates were high due to the monotonous and physically demanding nature of factory work. Furthermore, the skills required for these jobs were often simple and easily learned, making workers easily replaceable (Goldthorpe 2000). Consequently, these general skills did not provide workers with significant leverage over capital. The workers' dispensability resulted in low structural power, with factory owners maintaining dominant control over the conditions and compensation of labour.

2.2 Digital Revolution

Contrarily, the second wave, defined by the advent of computers and the Internet, presented a dichotomous effect. It had eroded the demand for certain specific skills related to routine tasks while simultaneously enhanced the necessity for general skills such as analytical thinking, problem-solving, and digital competency (Autor, Levy, and Murnane 2003). Jobs characterised by routine tasks, e.g., bookkeeping or telephone operators, were among the most affected during this wave. These roles, which were *specialised*[7] and demanded expertise in a particular field (data entry and telephone switchboard operation respectively), became replaced by computers and machinery designed to perform such tasks with greater speed and accuracy. Also, the Digital Revolution *transformed* the nature of many jobs and created entirely new ones, significantly increasing the demand for a different set of skills. Digital proficiency emerged as a vital *general* skill. The ability to operate computers, use software, and navigate

[7] NB: not only specialised and routine tasks got replaced; general and routine too: for example, certain assembly line operations were completely automated with robots.

the Internet became essential in a wide variety of job roles across different sectors (Bresnahan, Brynjolfsson, and Hitt 2002).

In addition to digital competency, general skills, such as analytical thinking and problem-solving, became overall more important in comparison with the times of Industrial Revolution given the increasing complexity and unpredictability of tasks in the knowledge-based economy. They facilitated the ability to adapt to changing technological environments and make informed decisions based on the analysis of a vast amount of data (Autor, Levy, and Murnane 2003). Furthermore, as Western economies transitioned to post-industrial/service sectors, businesses pivoted towards a customer-centric model with a heightened emphasis on human interaction. Consequently, Emotional Intelligence (EI) – the ability to understand and utilise one's emotions – has seen a surge in research interest (Mayer, Roberts, and Barsade 2008) highlighting the importance of general skills[8] in modern economies.

2.3 AI Revolution

The current wave of automation, characterised by the industrialisation of AI, echoes the dynamics of the second wave but on an amplified level. AI systems, compared to computers, have the capacity to undertake complex, *non-routine*, intelligent tasks (hence the name), thus encroaching on specific (i.e., narrow)[9] skill domains once considered exclusively human, namely diagnostic medicine or legal research (Brynjolfsson and McAfee 2014; Esteva et al.

[8] The EI is designated as a general skill, following the Becker's (1994) criterion of *transportability*, i.e., being able to carry it from one employer to another.

[9] **NB**: for the problem with naming, as well as a more substantial critique of the VoC approach to skills, see Streeck (2011). Especially illuminating is the idea of adding another dimension to skill taxonomy – market conditions, meaning the economic portability of the skill. Both astrophysicists and mathematicians have general skills but for the former "small and static [labour demand] … does not help them find employment [even in presence of broad skills]" (Streeck 2011, 18). Further studies may go further than this thesis in assessing the *impact of AI on skills with different economic portability.*

2019). Indeed, they can have a profound effect on *specific* skills as they tend to be occupationally unique (Streeck 2011). It is not clear, whether in the short-run AI would *replace* humans performing these tasks,[10] but for now, it is possible to conclude that AI can definitely augment and enhance human abilities. As evidenced by GitHub's (2023) AI-powered programming companion 'Co-pilot', or even the coding capabilities of ChatGPT, AI can indeed handle tasks like code-writing that are usually considered specific (i.e., narrow).

However, while AI is phasing out and/or augmenting some specific skills, it is also elevating the importance of *general* skills. For example, AI literacy is emerging as a crucial general skill. It comprises "a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace" (Long and Magerko 2020, 2). In essence, AI literacy is a skill that enables one to effectively operate AI systems in day-to-day and/or business operations. For instance, prompt engineering for ChatGPT can be considered an example of AI literacy. It is general (i.e., portable, and not tied to a particular industry) and helps to better tie the results to fit the preferences of the human operating it. In other words, while specific (narrow) skills such as coding, creative writing, or diagnostic medicine are being further exposed to AI, it becomes pivotal to have the skill to operate these AI solutions that augment/replace the human input.

At the same time, certain tasks are just not (yet) susceptible to automation. The studies of job polarisation show that both 'lovely jobs at the top and lousy ones at the bottom' (Goos and Manning 2007) are hard to automate. For 'lovely jobs' (as CEOs for example) rely on social connections and personal, *tacit*, knowledge that cannot be accounted for when training an AI

[10] I speak of tasks and not jobs as this study relies on a task-based approach to automation (Autor 2015) dividing the occupation into tasks (or competences) with varying prospects of automation, as would be evident from the methodological section.

system. On the other hand, 'lousy jobs' also rely on the low-pay *human* knowledge that is needed to support technological advancements, such as providing training data for machine learning systems like Amazon Mechanical Turk. Additionally, job tasks relying on dexterity, physique, and communication are hard to automate as robots are clunky (Autor and Dorn 2013). Hence, the jobs that are theoretically threatened the most are those that (not accounting for seniority or skills) do not rely on interpersonal communication or physical activities.

Overall, this section has sought to illuminate the future of work in the context of AI and automation, as well as their impacts on the skills landscape. Across three waves of automation – the Industrial Revolution, the Digital Revolution, and the AI Revolution, I identified changes in the in-demand skills as influenced by automation, reflecting the transformative nature of technological progress. The dichotomy of perspectives on automation underscores the complexity of this issue. Proponents advocate for automation's capacity to bolster productivity, create new job opportunities, and enhance the quality of life. In contrast, critics warn of the unprecedented nature of the current wave of automation, which could potentially render a wide swath of jobs obsolete and exacerbate social inequalities (as discussed in Section 1.1).

Across these automation waves, it is clear that each wave created and destroyed jobs but also transformed the nature of work and the skills required. The Industrial Revolution marked a shift from specific to general skills, largely devaluing artisanal expertise in favour of low-skilled labour that can be used as a 'machine operating another machine' on a production line. The Digital Revolution further decreased the demand for specific routine task-oriented skills while elevating the need for non-routine general skills like analytical thinking or digital competency. The AI revolution, still in progress, threatens to automate even highly specialised tasks, challenging the demand for specific skills once again while amplifying the importance of broad, adaptable skills and AI literacy. It is critical to note, however, that while AI and automation bring significant transformation, they do not signal the end of work. Certain tasks, particularly those involving complex human interactions or relying on human physique, remain resistant to automation. Lastly, it is important to consider that each wave of automation created a new job class – people who worked to maintain the machines that were used in production. Starting with line operators on a conveyor belt, progressing to engineers maintaining complex machinery during the Digital Revolution, and now encompassing a new class of 'digital engineers' – data specialists who maintain and process enormous amounts of data while developing product solutions – their jobs remain essential to the operation in every wave of automation.

2.4 Research Question & Hypothesis

As I have shown, the field of automation research has witnessed significant development, with various studies outlining different prospects for (and consequences of) automation. However, a critical analysis of the existing comparative political economy literature (in this case, VoC scholarship) reveals a gap. Although many studies have measured the potential for automation (Frey and Osborne 2013; Arntz, Gregory, and Zierahn 2016; Webb 2019), there *is a notable absence of the comparison of the propensity to automation of job tasks between LMEs and CMEs*. Consequently, my research seeks to address this gap by exploring the following research question:

RQ: Considering their distinct innovation patterns and skill systems, does the proclivity for job task automation diverge between Liberal Market Economies (LMEs) and Coordinated Market Economies (CMEs)?

Consequently, I hypothesise that:

The differences in the innovation patterns between the LMEs and CMEs would lead to a variation in their predispositions to automate job tasks with LMEs displaying a higher degree of potential automation due to their propensity for radical, disruptive innovation, which often challenges traditional production systems. Furthermore, LMEs tend to foster more competitive market conditions that encourage rapid technological advancements and risk-taking thus giving a stimulus for the actors to adopt cutting-edge technology (e.g., AI). Additionally, the distinct skill systems that these economies promote (general vs. specific) may influence the propensity for job task automation. As previously demonstrated, specific skills could, theoretically, be replaced by AI relatively soon, while general skills are gaining importance. However, in LMEs, these two mechanisms might offset each other. While AI poses a greater threat in LMEs due to their predisposition towards radical innovation and thus faster adaptation of AI, it is also less dangerous because it presents fewer risks to jobs requiring general skills. Considering that LMEs tend to prioritise general skills, AI might actually have an enabling effect rather than a replacing one in these countries. Lastly, another potential contributing factor to the differences in automation perspectives between LMEs and CMEs could be the geographical concentration of leading technology firms. Many of these firms are predominantly located in LMEs countries, such as the United States or the United Kingdom. This geographical advantage may accelerate the adoption of new technologies and automation processes in these economies.

Chapter 3. Methodology

Considering the complexity of the methodology employed and the numerous 'operational' decisions that had to be made, I have structured the discussion of methods as follows. Firstly, I provide an overview of the methods used, briefly discussing the data, methods themselves, and their respective limitations. However, for readers interested in a more comprehensive description of the *text analysis* techniques used as well as possible shortcomings stemming from data processing, subsequent sections delve into a detailed discussion of the technical decisions taken and their justifications.[11]

3.1 Overview

At its core, this study adopts the methodology proposed by Webb (2019). Following his approach, I extract the texts of titles of patents and job requirements.[12] Subsequently, from both job tasks and patents verb-noun pairs (VN) are extracted using the dependency parsing algorithm (see section 3.2). Then, the text undergoes pre-processing (lemmatisation, removal of stop-words, etc.). Next, VN patent pairs are aggregated using higher level semantic abstractions (Gangemi, Guarino, and Oltramari 2001) (see section 3.3). Then, the number of occurrences of the VN patent pair is divided by the total number of pairs to obtain their relative frequency. Then, each VN pair in the job description dataset is assigned the same frequency as the corresponding patent pair. By taking the average of the scores within one job (i.e., between tasks within one job), I derive an AI-exposure score that enables meaningful comparisons across different economies. The methodology is presented graphically in Figure 1.

[11] **NB**: the script and data used in this thesis have been **uploaded to the dedicated public GitHub repository** to comply with the Data Transparency regulations. The repository can be accessed via: https://github.com/nikitazerrnov/MA_thesis_CEU (see README.md for further details).

[12] Job requirements, tasks, and competences are considered synonyms for the purposes of this research.

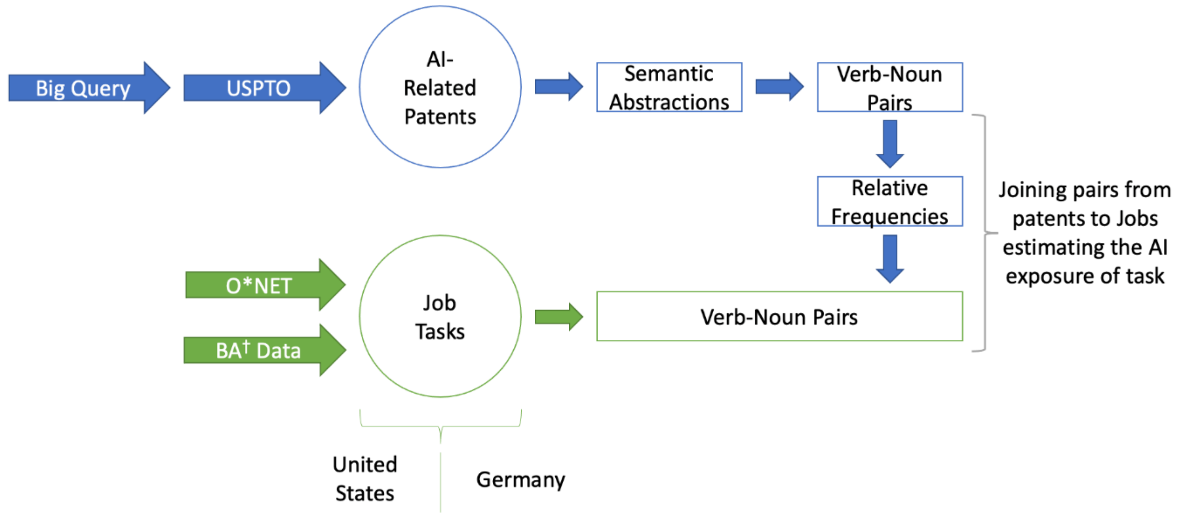

† Bundesagentür für Arbeit

Figure 1. Methodology

The choice of methodology is justified on several grounds: 1) the simple count of AI-related patents (akin to H&S) will not fully capture the complexity and nuances of AI's potential impact on job tasks; 2) departing from Frey and Osborne (2013), this study views a job as a collection of different tasks, each with various probability of automation, thus presenting a more accurate picture of automation. Additionally, I use computation and not expert reviews and imputation to measure the exposure score. Following H&S approach, I also focus on only two countries: the US and Germany to achieve the maximum possible divergence and thus a more illustrative comparison. The additional reason is data availability and required computational power: even processing data for these two countries requires working with terabytes of data and relying on cloud solutions (Google Cloud Platform) given the intensity of operations, on which I delve deeper in the following sections. However, while the data for the US labour market is available in a convenient format (Occupational Information Network (O*NET) database is maintained by the U.S. Department of Labor, Employment & Training

Administration), the data for Germany required extensive processing and establishment of concordance between O*NET and the German *Bundesagentur für Arbeit* database.

The chosen methodology is not devoid of limitations. Firstly, any automatic text-extraction method cannot comprehend the context and lacks 'human judgement', though the dependency parsing algorithm achieves > 90% accuracy in benchmarks (Honnibal and Johnson 2015). Secondly, the internal validity of the findings can be impaired given the diversity of the input data structures – it is nearly impossible to establish a perfect concordance between the US and German occupations; and the machine translation of the German tasks into English might have resulted in several mistakes and incorrect translations. Also, I processed only 1,000,000 patents (given the computational difficulties described below). While it is possible to implement stratified sampling, i.e., randomly sampling from each class (strata), when working with certain types of data, this approach is not viable when dealing with patent titles as they cannot be classified without advanced machine learning techniques. Therefore, when randomly sampling patents, there is a risk that data on some unique inventions, not seen in other patents, was lost. However, semantic abstractions mitigate this issue to an extent. Having outlined the methodology, I delve deeper into the technicalities of data processing.

3.2 US Job Tasks

As mentioned earlier, the data for the US is easily available. The Occupational Information Network (O*NET) database has a special 'Tasks Statements' dataset upon which further analysis is built. Overall, it contains information about 923 unique occupations (indexed by the Standard Occupational Classification (SOC) codes). There are approximately 20 job descriptions per title – with the minimum being 4 (for shampooers) and the maximum being

40 (for teachers). The job tasks range from those of CEOs that include, *inter alia*, “direct[-ing] or coordinat[-ing] an organization’s financial or budget activities to fund operations, maximize investments, or increase efficiency” to baristas[13] who are tasked with “weigh[-ing], grind[-ing], or pack[-ing] coffee beans for customers.” To start, task descriptions were pre-processed using tokenisation by breaking a string sequence (essentially, a text of a job task description) into individual units, aka ‘tokens’. For this analysis, I tokenised by words, though other forms of tokens, e.g., bigrams – a token consisting of two units – or n-grams can be processed (Jurafsky and Martin 2009).

After tokenisation, I employed a dependency parsing algorithm developed by Wijffels, Straka, and Straková (2023) on the basis of de Marneffes et al's. (2021) Universal Dependences framework for morphosyntactic annotation of text. Essentially, dependency parsing is a linguistic technique used to analyse the grammatical structure of a sentence based on the dependencies between its words. It creates a dependency tree, where the nodes represent the words in the sentence, and the directed edges between the nodes denote the grammatical relationships (dependencies) between the words/tokens. Each dependency relation expresses the type of grammatical relationship between the head and the dependent (modifier). The head of the sentence is usually a verb, and every other word in the sentence has a path leading to the head, which reflects the principle of ‘dependency’ and the head-initial nature of a language (Kübler, McDonald, and Nivre 2009). This is the main principle by which Verb-Noun pairs were identified. Since dependency parsing focuses on how words in the sentence relate to each other, it is important to preserve their original form. Therefore, dependency parsing is done *before* lemmatisation.[14]

[13] Baristas are a *sui generis* occupational category.
[14] Separate lemmatisation was done only for pre-processing purposes in the creation of Figures 3 and 7. Otherwise, the UDPipe package for R allows for both the identification of dependency relations and lemmatisation.

Figure 2 provides a snapshot of how this algorithm works. To illustrate the work of the algorithm, I used a sentence "*The child eagerly unwrapped the gift, revealing a shiny toy car*."[15] It is possible to break down this sentence in the following way: 1) articles are assigned a *det* POS (part-of-speech) – given they are attached to their nominal head – a noun ('the' relates to 'child' in the beginning), 2) 'child' is a nominal subject – '*nsubj*' an actor that performs an action described by the verb, 3) the adverb 'eagerly' modifies (hence, '*advmod*'), 4) the words 'shiny'[16] and 'toy' are *compounds* to modify the noun 'car'), 5) 'car' itself is an object for the verb 'revealing', and 6) lastly, the verb 'revealing' is used as an adverbial clause modifier, '*advcl*', to modify the verb 'unwrapped'.

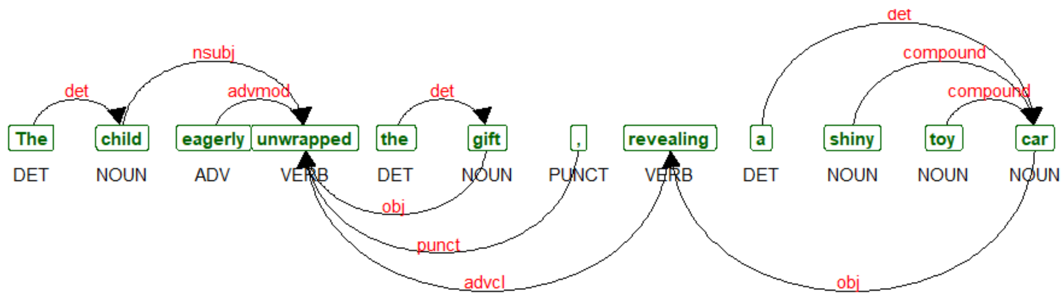

Figure 2. Working of the UDPipe dependency parsing algorithm: tokenisation, POS tagging & dependency relations

After establishing dependency relations between the words in every occupation, I extracted verb-noun (VN) pairs from the job tasks by *taking the main verb and its corresponding object*. Then, I lemmatised the VN pairs by reducing a word to its base or dictionary form – lemma. This technique differs from stemming, which simply removes (truncates) the endings of words. In essence, lemmatisation considers the morphological meaning of words. Thus, it is

[15] Note, that the performance and inner workings (treebanks) differ between the languages and even within different styles of one language.

[16] Interestingly, the algorithm is proven to be slightly faulty in this case, as it identifies the word 'shiny' as a noun while it is clearly an adjective. While it is not a serious problem, as dependency relation is unaffected by it, it might be the case that the algorithm was confused by the multiple compounds (De Marneffe and Manning 2008).

comparatively harder to perform as it requires a detailed dictionary and morphological analysis to correctly transform words to their base form. For example, if the dictionary is unaware of the homonymic nature of the word ‘leaves’ – it may be incorrectly lemmatised. In the sentence ‘He *leaves* for the university at 7am”, lemmatising ‘leaves’ to ‘leaf’ would be incorrect. At the same time, stemming would avoid such a problem by merely truncating ‘leaves’ to ‘leave’ which, in this case, is the right choice. However, stemming would not be able to substitute, for instance, possessive pronounces (e.g., ‘*His* boots were lying on the floor’) for personal pronounces. Lemmatisation was preferred as by reducing word inflections (different endings for plurals, comparatives, tenses, etc), it consolidates similar words, reduces data complexity, and improves the performance of dependency parsing algorithm (Balakrishnan and Ethel 2014).

After lemmatisation, I analysed the resulting VN task pairs. Overall, the data processing resulted in 29,755 VN pairs. However, given the textual nature of the data, I expected that a single job task description to contain several VN pairs (be in two sentences or in one compound sentence). Hence, to preserve this heterogeneity for data storage, a two-level hierarchical structure was implemented: the first, lower, level stores VN pairs within a single job task, while the second, upper, identifies the occupation. For example, for a barista, the data has 12 VN pairs, two of which come from a single job description. Overall, the pre-processing of the US job tasks was rather straightforward. Additionally, given that the same steps were applied in the German case – I would not describe them in details. Instead, I will focus on the limitations and difficulties. For now, though, I move to the discussion of the US patents.

3.3 AI-Related Patents

While the original methodology by Webb (2019) relied on the selection of patents based on the combination of key-words in their Cooperative Patent Classification (CPC) codes, I used a different approach for several reasons. Firstly, *human resources unavailability* – Webb employed a team of Research Assistants to generate all possible including/excluding key-words for identification of AI patents who are not available for this project. Secondly, *data availability* – since the day of Webb's publication United States Patent and Trademark Office (USPTO) released an "Artificial Intelligence Patent Dataset (AIPD)" (Giczy, Pairolero, and Toole 2022). Scholars used advanced Machine Learning (ML) techniques to identify US patents and pre-grant publications that contain one or several of the AI features, such as ML, Natural Language Processing (NLP) techniques, computer vision, etc. However, AIPD contains only a publication number of a patent, and not any other patent-related information (e.g., assignee information, patent *title*/text/abstract, date of granting, etc.). Hence, following Webb's approach, I relied on querying and subsetting the Google Patents Public Data (IFI CLAIMS Patent Services 2016) to retrieve the title of the patent. Google's repository contains information on 146,883,321 patents worldwide. To acquire a final dataset for analysis, I used patent numbers from AIPD to filter the Google Patents database for the AI-related US patents and their titles using SQL queries.[17] Overall, the final dataset contained information on 13 million AI-related patents.[18] However, due to computation limitations all 13 million cannot be

[17] I excluded patents that did not have a title. However, as they accounted for less than 1% of the data, removal should not skew the results.

[18] Importantly, since the patents were chosen by an ML binary classifier (i.e., predicting whether a patent is AI-related or not), this selection is subject to usual ML constraints of false-positives and false-negative predictions (Guido and Müller 2016). Given the nature of classification, I suppose it would be *costlier to miss an AI-related patent (false negative prediction)* than to misidentify non-AI as an AI-related (false positive). Thus, greater attention should be paid to the recall (sensitivity) metric, which is 0.38, meaning the model identifies only 38% of actual AI-related patents as AI-related. On the other hand, the model correctly predicts the class of the patent in 87% of the cases, meaning that it is probably good in identifying non-AI patents as non-AI. Notably, the true

pre-processed.[19] Hence, I randomly sampled 1 million patents and performed the parsing on the titles. This resulted in 8 million tokens with a mean number of words in a title being 6.

However, since it is possible that two or more patents can be closely related – i.e., describe similar inventions in the same sphere, it is better to aggregate them to achieve a more significant frequency of the VN pair extracted. Indeed, if no aggregation is performed, there would be a few pairs that appeared most often in the data, but all others would have a frequency of 1 since that particular word combination was not used. To avoid this problem, aggregation is needed, however, a typical approach proved unscalable. Usually, for small datasets fuzzy-matching with Levenstein Distance (LV) is used. For this to work, it would have required pair-wise calculations of LVs between all pairs of VN pairs. This quickly runs into $\Theta(n^2)$ problem (Gajentaan and Overmars 1995) – when the performance of the algorithm depends on the squared input. In plain words, it means that to calculate LVs between all pairs of VN pairs (of which there are 400,000) it would be necessary to process (400,000^2)/2 = 80 billion observations. Since it is clearly impossible, I relied on another approach – using the WordNet dictionary (Miller 1995; Fellbaum 2010). I created a higher-level abstraction for the nouns constituting a pair by taking a hypernym (i.e., a broader term) of a third level. For example, the word 'apple' would become 'food' through the chain 'apple → edible fruit → fruit'. This allowed to reduce the number of unique pairs (minimising the overall *noise* in the data) and thus increased the frequency of the aggregated pairs.

value for prediction was supplied by human examiners, that also do not identify them with 100% accuracy and this model still outperforms others (Giczy, Pairolero, and Toole 2022, tbl. 5).

[19] For example, in deployed Google Vertex AI Python notebook, only the process of dependency parsing would have taken 14 hours, which is not financially sustainable given Google's billing rates.

Overall, the pre-processing of the patent data revealed several of the limitations of the chosen methodology. Firstly, the data might not have contained as many *AI-related* patents as expected (see footnote 18) which might have influenced the analysis. Secondly, taking only a sample in a situation where it is impossible to evaluate its representativeness could have introduced a bias in the data, as some AI-related inventions might have got excluded. Lastly, the semantic abstractions, while a quick way to increase the relative frequency of pairs, might have distorted the original meanings and over-generalised the nouns, resulting in poorly-matched pairs. At the same time, these compromises were necessary as the volume of data was simply too large for an unabridged examination. Having discussed the peculiarities of analysis for the US, I advise moving to the German case.

3.4 German Job Tasks

As no suitable dataset for Germany was available, one was created from scratch. The steps taken to build it are as follows: 1) the concordance between the US and German occupations was established. Though different approaches exist,[20] I relied on manually 'agreeing' the jobs. For example, the job title corresponding to chief executive in the US is the position of *Geschäftsführer* in Germany; 2) then job descriptions were manually extracted from the database of *Bundesagentur für Arbeit* (Federal Employment Agency) of Germany. It contains detailed descriptions of jobs, together with the tasks one will perform on the job (*Aufgaben und Tätigkeiten* section). Overall, the data for 122 occupations is available[21] with, on average, 1600 characters per the job description (the US data had 1950); 3) the data was translated using

[20] Future studies may use the published "crosswalk between ESCO and O*NET" (European Commission 2022) to extract and compare job tasks in Europe with the US based on their ESCO and corresponding SOC code (for the US). This study did not employ this approach since the primary German data source (*BA*) is structured neither by ISO nor by ESCO codes.

[21] I attempted to extract more general occupations compared to the US data.

Google Translate API into English, and lastly, 4) it underwent the same pre-processing as the US (so it will not be described in details). This approach has two main limitations: firstly, *manual data collection*: while there is an API for the job search portal ('Arbeitsagentur Jobsuche API' [2021] 2023), I could not extract the relevant section with it, thus the job descriptions had to be manually pulled out which significantly decreased the throughput. Secondly, using *machine translation* was a sub-optimal choice (given the possibility of incorrect translation of grammar structures which could undermine the work of the dependency parsing algorithm), but it is scalable, which was a priority given the amount of data. Having explored the technicalities and trade-offs involved in data processing it is time to finally consider the results.

Chapter 4. Results & Discussion

To structure the discussion of results, I firstly overview the preliminary findings for the US (derived from pre-processing of the job tasks and patents), then analyse the automation prospects in the US and discuss its implications. Next, I repeat the same process for Germany and in conclusion, compare the two economies.

4.1 United States

To start, I analyse the distribution of lemmas of the US job tasks. Figure 3 presents the respective bar chart of the frequency of the top 20 lemmas (without considering Part-of-Speech tagging).

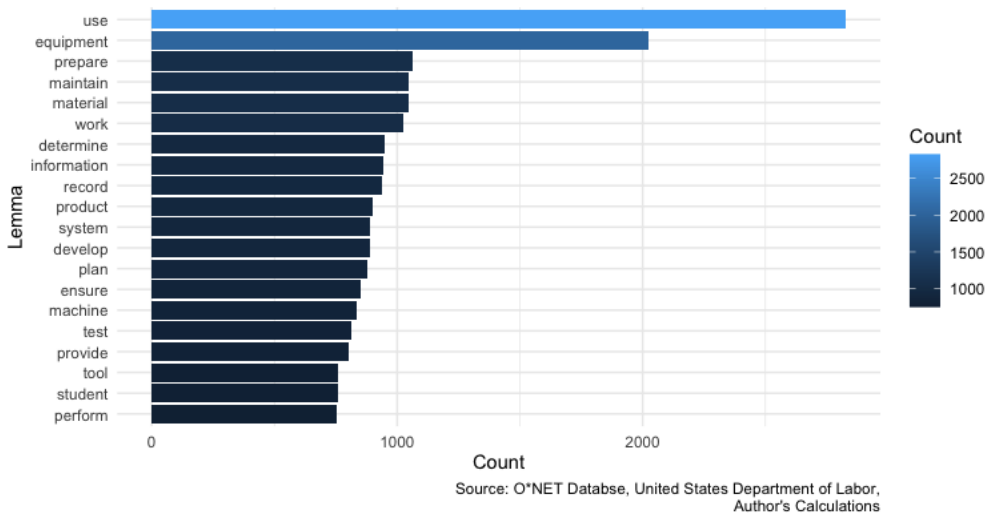

Figure 3. Frequency of Lemmas in US Job Tasks

The most frequent lemma found in the job tasks is ‘student’ and ‘research’ – possibly indicating the increasing importance of general skills. At the same time, workers are expected to ‘use’ different ‘equipment’, ‘material’, ‘tools’, etc. Interestingly, Zipf’s law – stating that the frequency of occurrence of a word is inversely proportional to its rank in frequency distribution (Piantadosi 2014), a widely held assumption in NLP – does not seem to hold with regard to job task descriptions. As can be observed, starting from the third lemma, the frequency seems to be relatively stable, while it can be expected that a third lemma would appear twice as often as the fourth one, but in reality, the difference is minimal. Otherwise, further insights can hardly be extracted from this. At the same time, more conclusions can be drawn from the analysis of the Verb-Noun pairs in the same data.

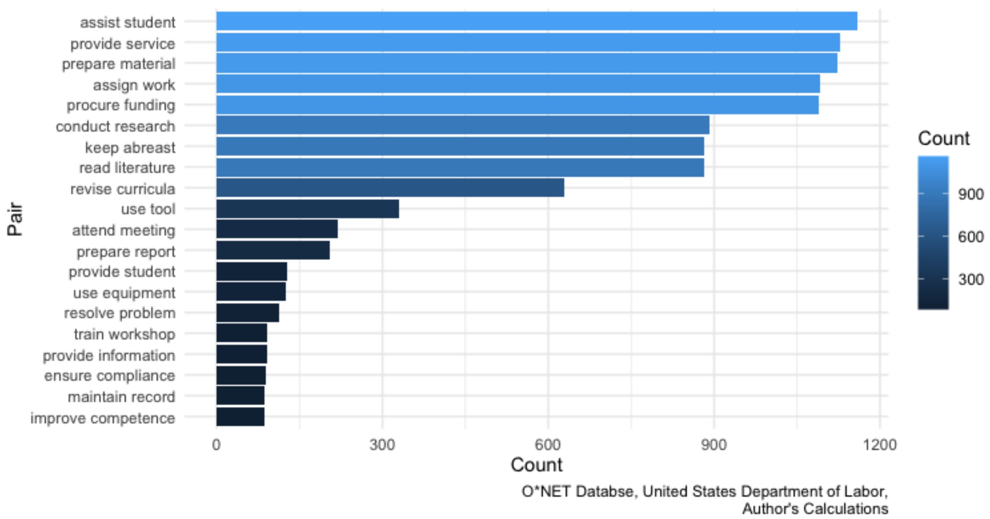

Figure 4. Most Common Verb-Noun Pairs in US Job Tasks

Figure 4 depicts the 20 most common Verb-Noun pairs in the whole job tasks dataset. To start, the second most frequent VN pair is ‘provide service,’ which might hint at the post-industrial character of the US economy, as most people are expected to offer some kind of services. Additionally, one can see a ‘procure funding’ pair in the top-5 highlighting a high degree of the financialisation of the US economy. As for the distribution, it is evident from Figure 4, that the frequency of pairs drops significantly at approximately the 10th pair signifying the non-uniform distribution of skills demanded from people on the labour market as employers may demand only certain skills across different industries and firms, which concurs with the scholarship (Cunningham and Villaseñor 2016). Also, Figure 4 may highlight the preference of the US market for general skills. With the exception of the VN pairs 'train workshop' and 'use equipment,' which indicate specific (i.e., narrow) skillsets related to using particular industrial equipment, and pairs related to students which are harder to interpret, the remaining pairs can be understood in terms of general skills. For instance, 'prepare material' and 'conduct research' reflect information-gathering academia-like abilities,[22] while "problem-solving" represents an essential general skill that has been relevant since the days of the Digital Revolution, as discussed above. These skills indeed are not tied to any specific employer and have broad applicability across various industries.

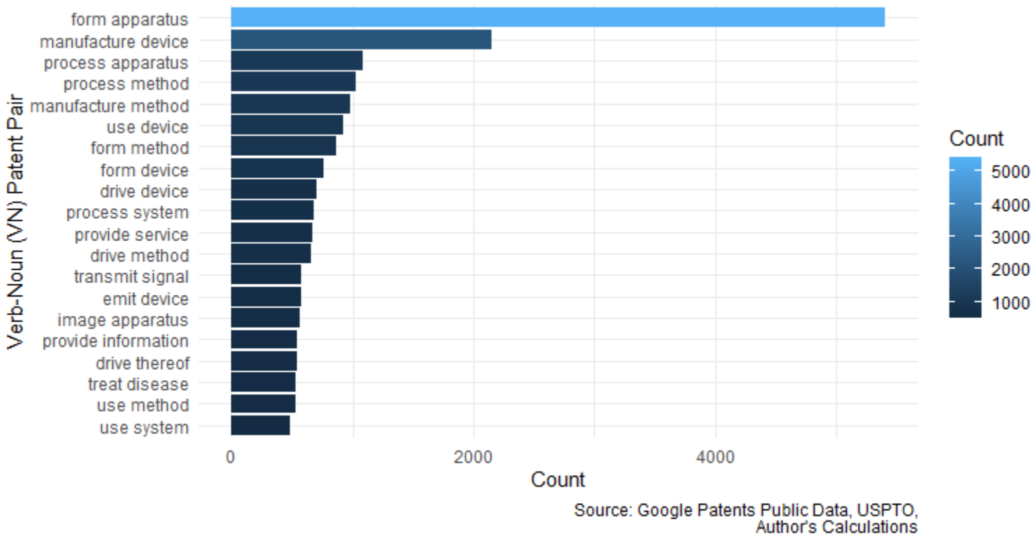

Figure 5. Most Common Verb-Noun Pairs in US Patents

As for the patents, as Figure 5 shows, the most common pair is 'form apparatus' which occurs almost five times as often as the others, followed by different use cases of 'devices' – 'manufacture', 'use', etc. This presents an issue, given that many VN pairs are essentially

[22] To delve into the disambiguation of the conceptualisation of general skills as broad and high in an academic context, see Streeck (2011).

(semantically) similar – they depict the same phenomenon (e.g., ‘use device’ and ‘use system’). However, this cannot be captured by the simple count of unique VN pairs. Hence, to remove this noise from the data, aggregation of VN pairs is indeed required (for details see section 3.3).

4.2 US Automation Prospects

Finally, Figure 6 shows the AI exposure score for the US. Across all the jobs, the mean exposure to AI automation is **0.02**%. Although the number cannot be compared with Webb (2019), as he did not provide a single mean automation index, and with other studies, given methodological differences, it can be used to compare the results with Germany. As for the US data, the index is rather skewed towards the lower end of the distribution, with most jobs hardly being automatable.[23]

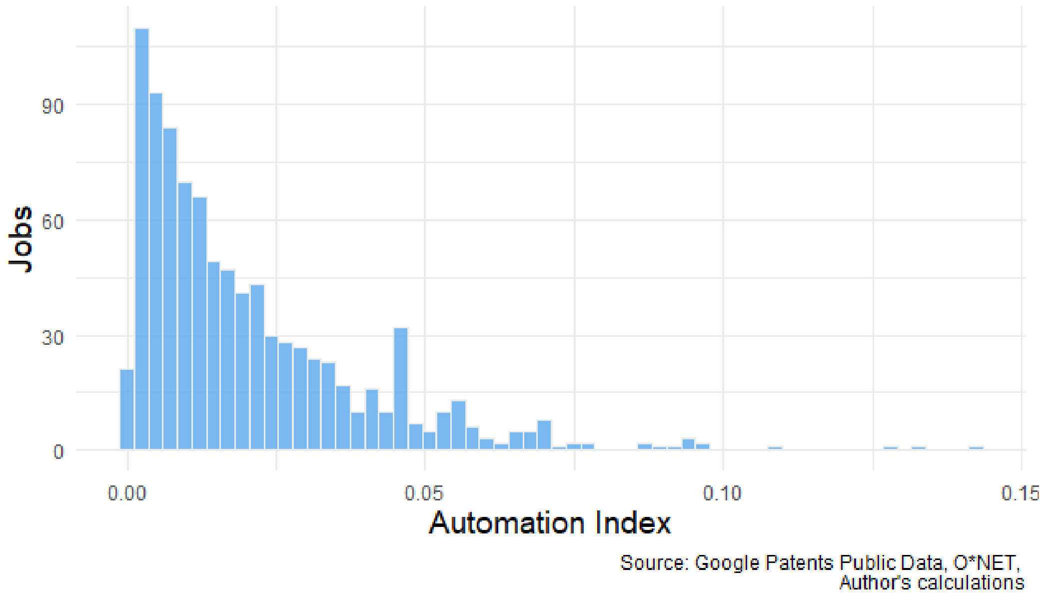

Figure 6. US Automation Prospects

[23] Here I speak about jobs and not tasks, as it is a more vivid comparison. Furthermore, given that a job may consist of several tasks with different prospects of automation, by comparing automation among the jobs I average the automation prospects of the tasks included.

As for the ‘tail’ of the distribution – the occupations which are most exposed to automation – the highest score is attained by physicists (0.14), followed by sound engineering technicians (0.13), pipelayers (0.12), fishing and hunting workers (0.10), crematory operators (0.09), entertainment and recreation managers (0.09), geographers, and low vision therapists (0.09). The overview of the most and least exposed jobs to automation can be found in Table 1. In general, it can be concluded that while most of the jobs are not directly threatened by AI, *jobs involving highly specific skills are threatened the most* (physicists and technicians). At the same time, as predicted by the theory, a lot of high-end/non-routine, communicative, and/or manual (i.e., relying on human movements) jobs are among the least exposed (waiters, singers).

Table 1. Jobs Most and Least Exposed to Automation

Most Exposed	Physicists, Sound Engineering Technicians, Urologists, Captains (Pilots of Water Vessels), Court Reporters, Dentists, Airline Pilots
Least Exposed	Cooks, Travel Guides, Musicians and Singers, Business Managers, Salespersons, Nannies, Waiters, Investment Fund Managers

Thus, the findings mostly agree with the scholarship. Thus, Webb (2019, 40) found that ‘clinical laboratory technicians’ and ‘chemical engineers’ are occupations most exposed to AI. As they are both high-end jobs with a specific skillset, it corroborates the findings of this study. The same can be said about Felten, Raj, and Seamans (2019) and their theory of white-collar jobs being most exposed to AI. However, some cases cannot be accounted for by the theories of automation that I explored earlier. For example, it is not clear why investment fund managers are less susceptible to automation. They rely on highly specific and proprietary knowledge and do not have tasks involving human interaction or movements. Indeed, already now there are AI solutions on the market that completely automate investment management offering investment forecasts and predictions such as Axyion.AI (2023) partially funded by UniCredit Group. Nevertheless, I believe this discrepancy mostly stems from the data problems described above and not from the flawed theoretical assumptions. For example, when analysing selected aggregated VN pairs from patent data, it becomes evident that some were made too general, like 'manufacture know-how' or 'provide work,' and thus may not have found an overlap in the job description data. Alternatively, they might have overlapped with the pairs that were not mentioned often, thus driving the exposure score down.

At the same time, a 99:1 percentile comparison, analysing the median annual salaries of the top 1% who are the most likely to be affected by automation against the rest of the workforce somewhat strengthens the prevailing theory of automation that suggests that AI wave particularly affects high-end jobs.[24] Indeed, the mean of the average annual income for the top 1% is $78,780 while for the rest of the labour force it is $65,730. To perform a statistical comparison, a one-sided Wilcoxon rank-sum test with continuity correction was used (since the data is not parametric, i.e., non-normally distributed, a t-test is discouraged). It suggests that the data does not provide strong evidence to reject the null hypothesis – that the distribution of median annual salaries is identical between the groups (W = 140,471, p-value = 0.051). However, it should be noted that the p-value was fairly close to the threshold, thus it can be established that this comparison provides a 'Straw-in-the-Wind' support for the theory: it somewhat affirms its relevance but does not eliminate the rival hypothesis (Collier 2011). In other words, it confirms the theory about a higher susceptibility of higher-end jobs to automation but does not disprove its alternative. Having examined the automation exposure in the US, it is time to analyse the German case and compare the results.

[24] The data is provided by the United States Department of Labour (2023).

4.3 Germany

To start, I examine the lemmas in the job task descriptions in Germany. Their frequencies are presented in Figure 7. Overall, two phenomena warrant further examination. Firstly, the word ‘research’ is an outlier with 209 mentions. The same picture was seen in the US with the words ‘student’ and ‘research’ dominating the count. While it is possible to argue that given the gap between the most common lemmas, it is worth removing ‘research’, I believe it is not the case since it is not an uninformative word as it does convey the actual meaning of the task one will perform. Secondly, the common outlier as well as other highly frequent words (e.g., ‘course’ in the US, and ‘project’ in Germany), point to the possible convergence between the job tasks in these economies, as it is hard to find semantic differences between the lemmas. Interestingly, Zipf’s law also does not seem to hold in the German case either, which could warrant further examination by NLP scholars.

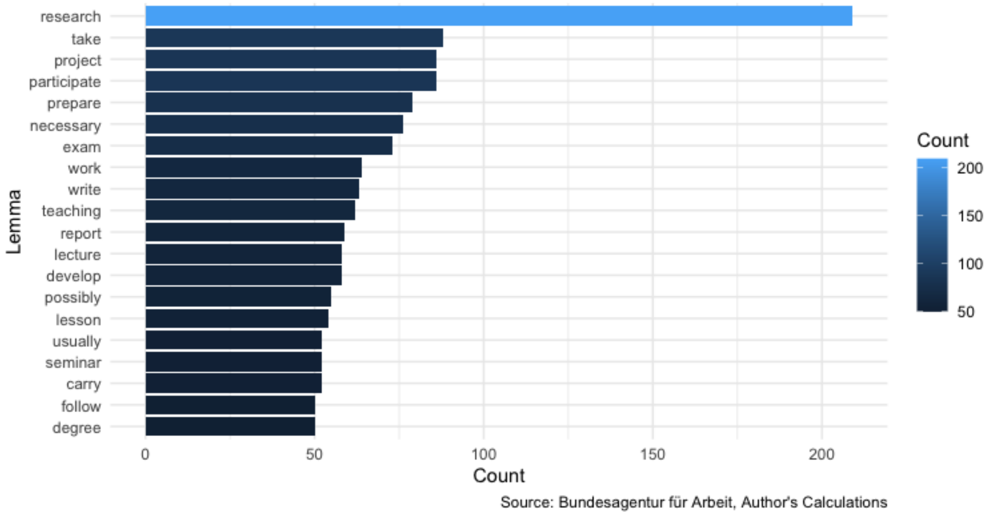

Figure 7. Frequency of Lemmas in German Job Tasks

The Verb-Noun pairs, however, present a more intriguing comparison. In Germany, three pairs emerge as relatively common: ‘take exam’, ‘prepare lecture’, and ‘hold course,’ as evidenced by Figure 8. Although the findings could potentially be influenced by the limited sample size (*n* = 657), it is noteworthy that the most frequent pair is ‘take exam’ - a combination not present in the US at all. This VN pair appears in the context of job descriptions, such as that for Agricultural Biology specialists: “*schriftliche Arbeiten korrigieren, Prüfungen abnehmen* [correct written work, take exams].” This finding serves to corroborate Streeck’s (2011) description of the German skill system, underscoring the importance of certified examinations and apprenticeships (*Berufsausbildung*) as key credentials of highly-skilled labourers (*Facharbeiter*) within the German context – which is alien to the Anglo-American system.

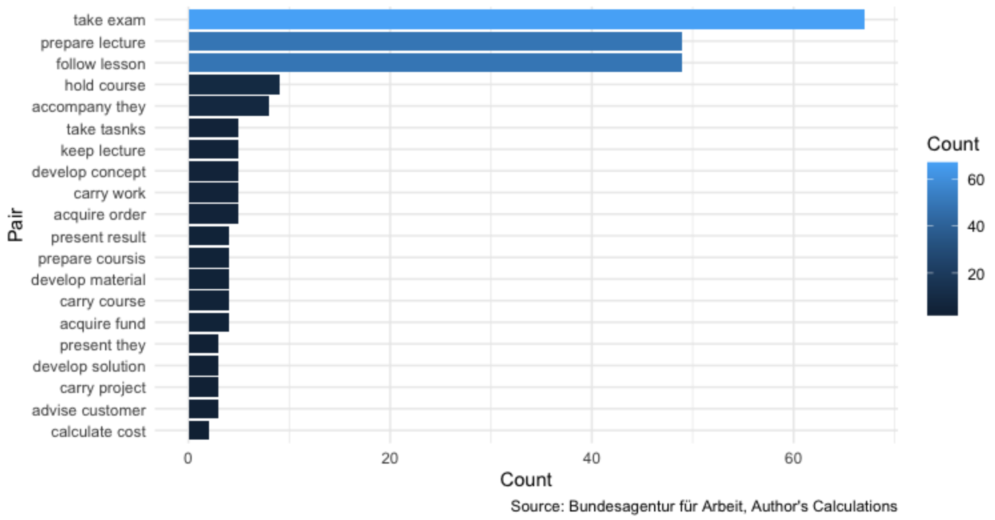

Figure 8. Most Common Verb-Noun Pairs in German Job Tasks

Other pairs, however, do not necessitate further investigation. They appear to be primarily centred around scholarly activities, echoing patterns observed in the US. This might underscore the growing significance of general skills in the era of AI automation and considering their ancillary role in CMEs, it alludes to potential disruptions these economies could face if they fail to adapt to evolving market realities. Moreover, both Germany and the US share the 'acquire fund' / 'procure funding' pair. While this fits well within the financialised US economy, it might also underscore a burgeoning demand for this skill in Germany, which is traditionally perceived as less profit-oriented, strengthening the argument about potential disruption.

4.4 Comparison of Automation Prospects

The overall picture for Germany differs from the US as expected by the literature, as depicted in Figure 9. The mean automation index across all jobs is merely **0.0051**% (compared to 0.02% for the US). At the same time, despite much lower prospects for automation, the pattern in the data distribution resembles that of the US. Most of the jobs are hardly automatable, but a small cluster of them has significantly higher prospects for automation (approaching 0.02 – the mean for the US).

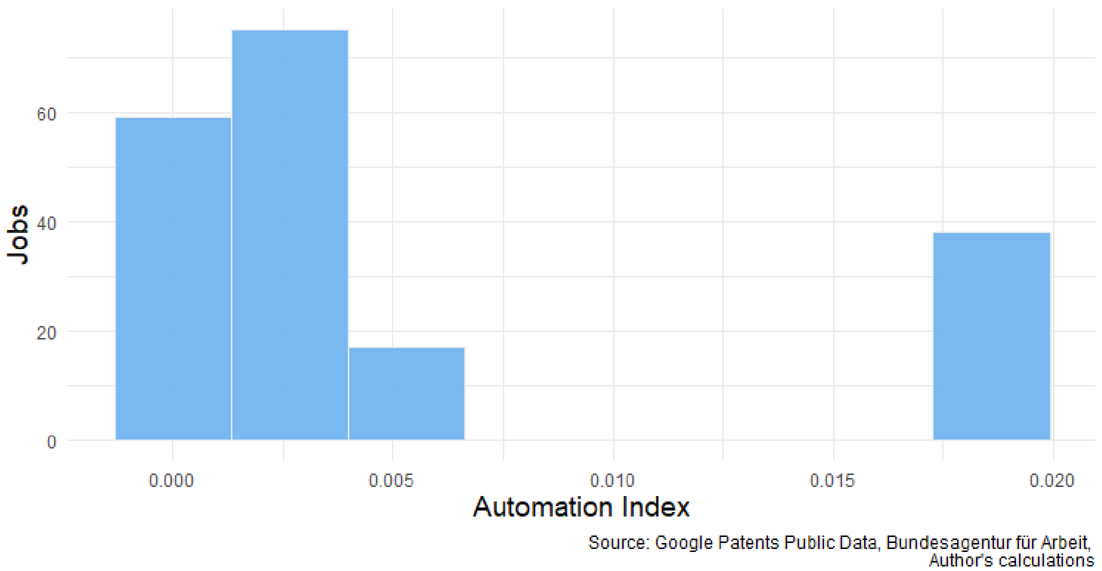

Figure 9. Automation Prospects for Germany

The similarities do not end there. The most/least exposed jobs are close to those in the States and the ones described by Webb (2019). The most exposed jobs are highly specific non-routine occupations, such as engineers, technicians, and doctors, while the least exposed are either routine (credit manager/food control officer) or those that require physique (fitness trainer), as can be seen from Table 2.

Table 2. Jobs Most and Least Exposed to Automation in Germany

Most Exposed	Construction Supervisor, IT-Technician, Epidemiologist, Construction Engineer
Least Exposed	Food Control Officer, Fitness trainer, Credit manager

Thus, the initial hypothesis that predicted *a difference in AI automation exposure between the US and Germany given their divergent skill systems and innovation patterns has been confirmed* by the analysis. Indeed, the difference between the AI exposure score between these countries approaches an order of magnitude. The difference is also *statistically* significant (Welsh two-sample t-test, t = 317.52, df = 17704, p-value < 2.2e-16), though given the small scale of measurement the significance may have been distorted. Moreover, the patterns of job susceptibility to automation are remarkably consistent across both countries, reaffirming existing theories about higher exposure of high-skill non-routine jobs or jobs with specific skills to AI automation (Lane and Saint-Martin 2021). Indeed, despite variations in the overall level of automation exposure, *it is evident that the same job categories tend to be exposed in LMEs and CMEs*.

Simultaneously, the findings underscore the importance of further refining the data processing and analytical techniques, in addition to the ones outlined throughout Chapter 3. For instance, consider that of the 13,874 German job task VN pairs identified, matches were found in the patents for only 4,380 pairs. This implies that nearly three-quarters of the potential data points were, in essence, discounted from the analysis. Future studies may find a better way to incorporate the unmatched pairs into the analysis to increase the reliability of findings. Furthermore, future qualitative studies may focus on strengthening the causal link between the automation exposure and skill regime/innovation pattern to be able to definitely attribute the difference to the theoretical assumptions with the help of, for example, process-tracing techniques that are capable of identifying causal patterns (Beach and Pedersen 2013).

Chapter 5. Conclusion

In 1965, Gordon Moore famously predicted that the density of transistors on integrated circuits would roughly double every two years, a prediction now known as Moore's Law. This law, while anticipated to plateau in the 2020s, symbolises the exponential surge in computational power witnessed over the past half-century. Indeed, the progression of technology has been nothing short of revolutionary, as were the changes it brought to *labour markets* with general skills getting importance and to *polities* with the Western countries shifting to post-industrial, service-oriented economies. However, the widespread adoption of AI carries the potential to exert the same, if not even more profound, influence on labour markets. Its emerging impact is already visible through the systems such as ChatGPT and CoPilot.

Driven by this momentum, this study set out to explore and compare the prospects of automation between Liberal Market Economies (LMEs) and Coordinated Market Economies (CMEs) based on the cases of the United States and Germany. It is based on the framework of Varieties of Capitalism (Hall and Soskice 2001), underlying the differences in skill systems and innovation patterns between the economies and theories of the technological impact on labour markets stipulating the distinct impact of the new wave of automation on jobs: non-routine jobs being affected more compared to the routine ones in the case of AI (Felten, Raj, and Seamans 2019; Lane and Saint-Martin 2021). The analysis was performed using quantitative text analysis methods that relied on the extraction of job tasks for the US and German occupations and the calculation of the relative frequencies of Verb-Noun pairs for these tasks based on the occurrence of the same pairs in the titles of AI-related USPTO patents to construct an index of exposure to AI of jobs.

This methodology, however, is not devoid of limitations. Thus, future studies may utilise either more data (patents and their texts or job descriptions in other countries) or more in-depth analysis (measuring the exposure to automation by technological sectors, using qualitative research techniques that excel at 'human judgement' in detecting meanings or using process-tracing to establish the causal link between the firms' operational patterns (VoC) and the potential exposure of jobs to automation). Additionally, while it is theoretically expected that certain jobs, like doctors, would *not* exhibit a high automation exposure (given the large inter-personal component), future studies may delve deeper into why this analysis seemingly contradicts this theory (I humbly believe the high exposure is caused by the absence of match of job tasks VN pairs that have a communicative component to the patent ones, thus biasing the overall index).

Notwithstanding the technical limitations, the initial hypothesis of the study was confirmed: I showed that the mean exposure to automation index across all jobs significantly differs between the two countries, standing at *0.02% for the US and 0.0051% for Germany*. Despite this variance, the patterns in data distribution were strikingly similar. The majority of jobs in both countries demonstrated *low prospects for automation, with a small cluster exhibiting significantly higher automation potential.* A closer examination of the occupations most and least exposed to automation revealed another convergence. Non-routine occupations with highly specific (i.e., narrow) skills, such as engineers, technicians, doctors, pilots, and physicists were most exposed to automation, while routine roles or those relying on human physique or commutations were among the least exposed. These findings lend further weight to theories predicting that high and/or specific-skilled jobs are mostly exposed to AI automation.

Furthermore, this study is pioneering in demonstrating that this trend is observable in both LMEs and CMEs. This may suggest that LMEs are comparatively better prepared for the impending wave of AI-driven automation. While the exposure patterns are similar, the workforce in LMEs appears to rely more heavily on the less automatable general skills, at least theoretically. Therefore, automation could have a more 'enabling' effect (Frey, 2019) in these economies, signifying a lesser disruption to the status quo. At the same time, I must reiterate that the variable of interest was *exposure* to automation, not threat or substitution potential. What it showed is the percentage of a job being exposed to the sa*me tasks that were described in the AI-related patents*, meaning these tasks theoretically can (but not necessarily will) be performed by a machine.

To conclude, I want to stress the societal and political ramifications of the study. While, for now, only a small share of jobs is even *exposed* to automation, the situation can drastically change in the coming future. Hence, it is imperative to develop policies and strategies to ensure a *just* transition towards an increasingly automated economy. This may involve fostering active labour market policies to strengthen *education and training systems* to equip more workers with the skills necessary in the modern economy and/or providing *social protection* for those potentially affected to lessen the negative consequences of automation.

Reference List

Aguiar, Mark, and Erik Hurst. 2007. 'Measuring Trends in Leisure: The Allocation of Time over Five Decades'. *The Quarterly Journal of Economics* 122 (3): 969–1006.

Akkermans, Dirk, Carolina Castaldi, and Bart Los. 2009. 'Do "Liberal Market Economies" Really Innovate More Radically than "Coordinated Market Economies"?: Hall and Soskice Reconsidered'. *Research Policy* 38 (1): 181–91. https://doi.org/10.1016/j.respol.2008.10.002.

'Arbeitsagentur Jobsuche API'. (2021) 2023. Python. Bundesstelle für Open Data. https://github.com/bundesAPI/jobsuche-api.

Arendt, Hannah, and Margaret Canovan. 2010. *The Human Condition*. 2. ed., [Nachdr.]. Chicago: Univ. of Chicago Press.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2016. 'The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis'. OECD Social, Employment and Migration Working Papers 189. Vol. 189. OECD Social, Employment and Migration Working Papers. https://doi.org/10.1787/5jlz9h56dvq7-en.

Autor, D. H., F. Levy, and R. J. Murnane. 2003. 'The Skill Content of Recent Technological Change: An Empirical Exploration'. *The Quarterly Journal of Economics* 118 (4): 1279–1333. https://doi.org/10.1162/003355303322552801.

Autor, David H. 2015. 'Why Are There Still So Many Jobs? The History and Future of Workplace Automation'. *Journal of Economic Perspectives* 29 (3): 3–30. https://doi.org/10.1257/jep.29.3.3.

Autor, David H, and David Dorn. 2013. 'The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market'. *American Economic Review* 103 (5): 1553–97. https://doi.org/10.1257/aer.103.5.1553.

Axyon.AI. 2023. 'Axyon IRIS® - Machine Learning for Investing'. *Axyon*. 2023. https://axyon.ai/axyon-iris/.

Balakrishnan, Vimala, and Lloyd-Yemoh Ethel. 2014. 'Stemming and Lemmatization: A Comparison of Retrieval Performances'. *Lecture Notes on Software Engineering* 2 (3): 262–67. https://doi.org/10.7763/LNSE.2014.V2.134.

Beach, Derek, and Rasmus Brun Pedersen. 2013. *Process-Tracing Methods: Foundations and Guidelines*. Ann Arbor: University of Michigan Press.

Becker, Gary S. 1994. 'Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education, Third Edition'. The University of Chicago Press. https://www.nber.org/books-and-chapters/human-capital-theoretical-and-empirical-analysis-special-reference-education-third-edition.

Black, Bernard S, and Ronald J Gilson. 1998. 'Venture Capital and the Structure of Capital Markets: Banks versus Stock Markets'. *Journal of Financial Economics* 47 (3): 243–77. https://doi.org/10.1016/S0304-405X(97)00045-7.

Braverman, Harry. 1998. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. 25th Anniversary ed. edition. New York: Monthly Review Press.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. 'Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence'. *The Quarterly Journal of Economics* 117 (1): 339–76.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. First Edition. New York London: W. W. Norton & Company.

Casper, Steven. 2001. ‘The Legal Framework for Corporate Governance: The Influence of Contract Law on Company Strategies in Germany and the United States’. In *Varieties of Capitalism*, edited by Peter A. Hall and David Soskice, 1st ed., 387–416. Oxford University PressOxford. https://doi.org/10.1093/0199247757.003.0012.

Collier, David. 2011. ‘Understanding Process Tracing’. *PS: Political Science & Politics* 44 (4): 823–30. https://doi.org/10.1017/S1049096511001429.

Cunningham, Wendy V., and Paula Villaseñor. 2016. ‘Employer Voices, Employer Demands, and Implications for Public Skills Development Policy Connecting the Labor and Education Sectors’. *The World Bank Research Observer* 31 (1): 102–34. https://doi.org/10.1093/wbro/lkv019.

De Marneffe, Marie-Catherine, and Christopher D. Manning. 2008. ‘Stanford Typed Dependencies Manual’. Technical report, Stanford University.

Dehesa, Guillermo de la. 2006. *Winners and Losers in Globalization*. Oxford: Blackwell.

Esping-Andersen, Gøsta. 1990. *The Three Worlds of Welfare Capitalism*. Princeton (N.J.): Princeton university press.

Esteva, Andre, Alexandre Robicquet, Bharath Ramsundar, Volodymyr Kuleshov, Mark DePristo, Katherine Chou, Claire Cui, Greg Corrado, Sebastian Thrun, and Jeff Dean. 2019. ‘A Guide to Deep Learning in Healthcare’. *Nature Medicine* 25 (1): 24–29. https://doi.org/10.1038/s41591-018-0316-z.

Estevez-Abe, Margarita, Torben Iversen, and David Soskice. 2001. ‘Social Protection and the Formation of Skills: A Reinterpretation of the Welfare State’. In *Varieties of Capitalism*, edited by Peter A. Hall and David Soskice, 1st ed., 145–83. Oxford University PressOxford. https://doi.org/10.1093/0199247757.003.0004.

European Commission. 2022. ‘The Crosswalk between ESCO and O*NET (Technical Report)’. https://esco.ec.europa.eu/en/about-esco/publications/publication/crosswalk-between-esco-and-onet-technical-report.

Fellbaum, Christiane. 2010. ‘WordNet’. In *Theory and Applications of Ontology: Computer Applications*, edited by Roberto Poli, Michael Healy, and Achilles Kameas, 231–43. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-90-481-8847-5_10.

Felten, Edward W., Manav Raj, and Robert Seamans. 2019. ‘The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization’. SSRN Scholarly Paper. Rochester, NY. https://doi.org/10.2139/ssrn.3368605.

Frey, Carl Benedikt. 2019. ‘The Descent of the Middle Class’. In *The Technology Trap: Capital, Labor, and Power in the Age of Automation*, 227–48. Princeton University Press. https://doi.org/10.2307/j.ctvc77cz1.

Frey, Carl Benedikt, and Michael A. Osborne. ‘The Future of Employment: How Susceptible Are Jobs to Computerisation?’, 114:254–80. University of Oxford: University of Oxford, 2013. https://doi.org/10.1016/j.techfore.2016.08.019.

Gajentaan, Anka, and Mark H Overmars. 1995. ‘On a Class of O(N2) Problems in Computational Geometry’. *Computational Geometry* 5 (3): 165–85. https://doi.org/10.1016/0925-7721(95)00022-2.

Gallup. 2018. *Gallup’s Perspective on The Gig Economy and Alternative Work Arrangements*. The Gallup Building, 901 F Street, NW, Washington, D.C., 20004.

Gangemi, Aldo, Nicola Guarino, and Alessandro Oltramari. 2001. ‘Conceptual Analysis of Lexical Taxonomies: The Case of WordNet Top-Level’. *Formal Ontology in Information Systems: Collected Papers from the Second International Conference*, October. https://doi.org/10.1145/505168.505195.

Gheaus, Anca, and Lisa Herzog. 2016. ‘The Goods of Work (Other Than Money!)’. *Journal of Social Philosophy* 47 (1): 70–89. https://doi.org/10.1111/josp.12140.

Giczy, Alexander V., Nicholas A. Pairolero, and Andrew A. Toole. 2022. 'Identifying Artificial Intelligence (AI) Invention: A Novel AI Patent Dataset'. *The Journal of Technology Transfer* 47 (2): 476–505. https://doi.org/10.1007/s10961-021-09900-2.

GitHub. 2023. 'GitHub Copilot · Your AI Pair Programmer'. GitHub. 2023. https://github.com/features/copilot.

Goldthorpe, John H. 2000. *On Sociology: Numbers, Narratives, and the Integration of Research and Theory*. 1st edition. Oxford UK; New York: Oxford University Press.

Goldthorpe, John H., and David Lockwood. 1963. 'Affluence and the British Class Structure'. *The Sociological Review* 11 (2): 133–63. https://doi.org/10.1111/j.1467-954X.1963.tb01230.x.

Goos, Maarten, and Alan Manning. 2007. 'Lousy and Lovely Jobs: The Rising Polarization of Work in Britain'. *The Review of Economics and Statistics* 89 (1): 118–33.

Graeber, David. 2019. *Bullshit Jobs*. New York London Toronto Sydney New Delhi: Simon & Schuster Paperbacks.

Guido, Sarah, and Andreas C. Müller. 2016. *Introduction to Machine Learning with Python: A Guide for Data Scientists*. Sebastopol, CA: O'Reilly Media.

Hacker, Jacob S., and Paul Pierson. 2010. 'Winner-Take-All Politics: Public Policy, Political Organization, and the Precipitous Rise of Top Incomes in the United States'. *Politics & Society* 38 (2): 152–204. https://doi.org/10.1177/0032329210365042.

Hall, Peter A., and David W. Soskice, eds. 2001. *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*. Oxford [England]; New York: Oxford University Press.

Held, David. 2006. *Models of Democracy*. Stanford University Press.

Honnibal, Matthew, and Mark Johnson. 2015. 'An Improved Non-Monotonic Transition System for Dependency Parsing'. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 1373–78. Lisbon, Portugal: Association for Computational Linguistics. https://doi.org/10.18653/v1/D15-1162.

Howarth, David, and Lucia Quaglia. 2014. 'The Steep Road to European Banking Union: Constructing the Single Resolution Mechanism'. *JCMS: Journal of Common Market Studies* 52 (S1): 125–40. https://doi.org/10.1111/jcms.12178.

IFI CLAIMS Patent Services. 2016. 'Google Patents Public Data'. BigQuery. https://bigquery.cloud.google.com/dataset/patents-public-data:patents.

Im, Zhen Jie, Nonna Mayer, Bruno Palier, and Jan Rovny. 2019. 'The "Losers of Automation": A Reservoir of Votes for the Radical Right?' *Research & Politics* 6 (1): 205316801882239. https://doi.org/10.1177/2053168018822395.

Jeng, Leslie A., and Philippe C. Wells. 2000. 'The Determinants of Venture Capital Funding: Evidence across Countries'. *Journal of Corporate Finance* 6 (3): 241–89.

Jurafsky, Dan, and James H. Martin. 2009. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. 2nd ed. Prentice Hall Series in Artificial Intelligence. Upper Saddle River, N.J: Pearson Prentice Hall.

Kenworthy, Lane. 2001. 'Wage-Setting Measures: A Survey and Assessment'. *World Politics* 54 (1): 57–98. https://doi.org/10.1353/wp.2001.0023.

Kitschelt, Herbert, and Philipp Rehm. 2014. 'Occupations as a Site of Political Preference Formation'. *Comparative Political Studies* 47 (12): 1670–1706. https://doi.org/10.1177/0010414013516066.

Kohn, Melvin L. 2001. 'Job Complexity and Adult Personality'. In *Social Stratification, Class, Race, and Gender in Sociological Perspective, Second Edition*. Routledge.

Kübler, Sandra, Ryan McDonald, and Joakim Nivre. 2009. 'Dependency Parsing'. *Synthesis Lectures on Human Language Technologies* 1 (1): 1–127.

Lane, Marguerita, and Anne Saint-Martin. 2021. 'The Impact of Artificial Intelligence on the Labour Market: What Do We Know So Far?' OECD Social, Employment and Migration Working Papers 256. Vol. 256. OECD Social, Employment and Migration Working Papers. https://doi.org/10.1787/7c895724-en.

Leslie, Stuart W. 1993. *The Cold War and American Science: The Military-Industrial-Academic Complex at MIT and Stanford*. New York: Columbia Univ. Press.

Long, Duri, and Brian Magerko. 2020. 'What Is AI Literacy? Competencies and Design Considerations'. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16. CHI '20. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3313831.3376727.

Marneffe, Marie-Catherine de, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. 'Universal Dependencies'. *Computational Linguistics* 47 (2): 255–308. https://doi.org/10.1162/coli_a_00402.

Mayer, John D., Richard D. Roberts, and Sigal G. Barsade. 2008. 'Human Abilities: Emotional Intelligence'. *Annual Review of Psychology* 59 (1): 507–36. https://doi.org/10.1146/annurev.psych.59.103006.093646.

Miller, George A. 1995. 'WordNet: A Lexical Database for English'. *Communications of the ACM* 38 (11): 39–41. https://doi.org/10.1145/219717.219748.

Mokyr, Joel. 1992. 'The Years of Miracles: The Industrial Revolution 1750–1830'. In *The Lever of Riches: Technological Creativity and Economic Progress*, edited by Joel Mokyr, 81–112. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195074772.003.0005.

Mokyr, Joel, Chris Vickers, and Nicolas L. Ziebarth. 2015. 'The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different?' *Journal of Economic Perspectives* 29 (3): 31–50. https://doi.org/10.1257/jep.29.3.31.

Pecchi, Lorenzo, and Gustavo Piga, eds. 2008. *Revisiting Keynes: Economic Possibilities for Our Grandchildren*. Cambridge, Mass: MIT Press.

Pethokoukis, James. 2017. 'Will Technology Enable Workers or Replace Them? A Long-Read Q&A with Daron Acemoglu'. *American Enterprise Institute - AEI* (blog). 2 October 2017. https://www.aei.org/economics/will-technology-enable-workers-or-replace-them-a-long-read-qa-with-daron-acemoglu/.

Piantadosi, Steven T. 2014. 'Zipf's Word Frequency Law in Natural Language: A Critical Review and Future Directions'. *Psychonomic Bulletin & Review* 21 (5): 1112–30. https://doi.org/10.3758/s13423-014-0585-6.

Pierson, Paul. 2002. 'Coping with Permanent Austerity: Welfare State Restructuring in Affluent Democracies'. *Revue Française de Sociologie* 43 (2): 369–406. https://doi.org/10.2307/3322510.

Robert Skidelsky. 2023. 'The Future of Work'. Presented at *The Future of Work - Humans and Machines Lecture Series by Robert Skidelsky*, Central European University, Quellenstrasse 51, Vienna, April 19. https://events.ceu.edu/2023-04-19/future-work-humans-and-machines-lecture-series-robert-skidelsky.

Shonfield, Andrew. 1969. *Modern Capitalism: The Changing Balance of Public and Private Power*. London: Oxford University Press.

Soskice, David. 1997. 'German Technology Policy, Innovation, and National Institutional Frameworks'. *Industry and Innovation* 4 (1): 75–96. https://doi.org/10.1080/13662719700000005.

Streeck, Wolfgang. 2011. 'Skills and Politics: General and Specific'. SSRN Scholarly Paper. Rochester, NY. https://doi.org/10.2139/ssrn.1781042.

Taylor, Mark Zachary. 2004. 'Empirical Evidence against Varieties of Capitalism's Theory of Technological Innovation'. *International Organization* 58 (3): 601–31.

Trajtenberg, Manuel. 1990. ‘A Penny for Your Quotes: Patent Citations and the Value of Innovations’. *The RAND Journal of Economics* 21 (1): 172–87. https://doi.org/10.2307/2555502.

United States Department of Labour. 2023. ‘National Occupational Employment and Wage Estimates’. https://www.bls.gov/oes/current/oes_nat.htm.

Webb, Michael. 2019. ‘The Impact of Artificial Intelligence on the Labor Market’. *Available at SSRN 3482150*.

Wijffels, Jan, Milan Straka, and Jana Straková. 2023. ‘Udpipe: Tokenization, Parts of Speech Tagging, Lemmatization and Dependency Parsing with the “UDPipe” “NLP” Toolkit’. https://cran.r-project.org/web/packages/udpipe/index.html.

Womack, James P. 2007. *The Machine That Changed the World: The Story of Lean Production*. Reprint edition. London: Free Press.

Zarifhonarvar, Ali. 2023. ‘Economics of ChatGPT: A Labor Market View on the Occupational Impact of Artificial Intelligence’. SSRN Scholarly Paper. Rochester, NY. https://doi.org/10.2139/ssrn.4350925.