Essays on Wealth and Human Capital Inequality

Thesis by

Max Schroeder

Submitted
in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Economics

Adam Smith Business School, College of Social Science
University of Glasgow
March 31, 2023

ABSTRACT

This thesis studies wealth and human capital inequality, with an emphasis on the interaction of the two inequalities with each other and with large changes in the aggregate environment of the economy, such as technological change or large economy-wide shocks.

The first chapter proposes a model to study the effect of differentiated, cognitive skill-biased, technological change (CBTC) on income and wealth inequality. Recent studies have documented rising income inequality over the past decades and linked this phenomenon to the increasing adoption of computerised equipment and technologies. Other work relates changes in idiosyncratic risk faced by households to changes in the distribution of wealth. In this chapter I try and connect both literatures, to assess the effect of technological change on wealth inequality mediated by changes to income risk. For this, I combine elements of the "Task-Skill" framework with a heterogenous agent incomplete market model. The model includes a production environment that accounts for differentiated skill demand on the firms' side and multidimensional skill supply on the worker's side. Using measures of cognitive and non-cognitive skills from a comprehensive panel dataset for the UK (Understanding Society), I calibrate the model with appropriate micro-estimates. The calibrated model captures relevant features of the income process and wealth distribution observed in the data. In contrast to the standard approach of approximating income risk through an AR(1) process, the model allows income risk to respond appropriately to changes in the demand for cognitive skills. I then use the model to assess the impact of cognitive skill-biased technological change in the UK over the period 1980 - 2016. The model suggests that CBTC can account for the bulk of increases in labour income inequality observed over that period and is generally consistent with stylized facts about changes to wealth inequality.

The second chapter studies the complex interactions that exist between the distrib-

utions of wealth and human capital amongst working-age individuals. In this chapter, I develop a general equilibrium incomplete market model with endogenous wealth and human capital to analyse the interactions between these two factors. Workers choose between investing in a safe asset or augmenting their stock of knowledge and skills (human capital), which makes them more productive in the labour market. Human capital, however, is risky since it is not equally valuable in every employment situation. I calibrate the model to the UK economy in the pre-Covid-19 period and analyse the interaction of wealth and human capital in the stationary equilibrium. I find that there are important non-linearities in human capital investments, with workers with low levels of wealth investing considerably less in accumulating human capital than their counterparts with more wealth. I then analyse the economic dynamics of the distribution of human capital in the aftermath of an unexpected economic shock, showing that wealth poorer households are more exposed to these shocks, implying that the distribution of wealth matters for the recovery of the economy following recessions. These results provide a potential explanation for the persistently low productivity observed after the 2008 recession in the UK, highlighting the role of worsening wealth inequality as generating endogenous barriers for lower wealth individuals to improve their working skills. Finally, I assess the impact of the Covid-19 pandemic and associated support measures in the UK. The model predicts that the UK economy will likely suffer a significant reduction in human capital in the aftermath of the Covid-19 pandemic, but targeted policy action has helped to reduce the impact of the crisis particularly for low-wealth households.

The third chapter studies the change in the multidimensional skill supply of university graduates in the UK. University graduates have highly differentiated skills, both compared to the general population and other graduates. Differences arise from differences in background, course of study and individual aptitudes and interests. In this chapter, I study the distribution of these different skills, investigating what types of skills graduates have, and how these vary between and within broadly defined subject groups as well as across time. To this end, I develop a model of occupational choice and wage determination for university graduates in the UK. Graduates differ with regard to two types of general skills: mathematical/technical and verbal/organisatorial, which are used with different intensities by different occupations. I structurally estimate the model to find evidence of changing multivariate skill distributions over time. I find that between 1994 and 2019, the typical graduates' level of mathematical skills increased by 140% while verbal skills decreased by close to a third. Looking closer at 5 different major subject categories, I find that this trend is driven by increasing specialisation for STEM and Business & Economics degrees and increasing generalisation among Arts & Humanities and Other Subjects. For most graduates mathematical/technical skills have become the single biggest contributing factor to their earnings, making up around 50% of their hourly wage compared to 27% in the mid-90s. Counterfactual simulations suggest that in the absence of changes to the subject-specific skill distributions, mean wages would be up to 8% lower, while wage inequality would be up to 5% larger. The results suggest that graduate skill supply has adjusted to changing labour market requirements.

ALTERNATIVE THESIS FORMAT

This is an alternative format PhD Thesis and includes three papers. The papers are presented in the following order:

1. Cognitive Skill-Biased Technological Change, Income & Wealth Inequality in the UK
2. The Interplay between Wealth and Human Capital Inequality – Implications for the UK's post-Covid-19 recovery
3. To what degree? Recovering changes in the UK's graduate skill distribution

I declare that this thesis has not been submitted for any other degree at the University of Glasgow or any other institution.

Max Schroeder
Glasgow, March 31, 2023.

This page is intentionally left blank.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

For my parents, who taught me to live curiously.

ACKNOWLEDGMENTS

I would like to thank my advisors Jim Malley and Konstantinos Angelopoulos who helped me at every stage of this near five-year journey. Their feedback, advice and support has been invaluable and will always be appreciated. I would also like to thank Rebecca Mancy and the rest of the COVID-19 project team for providing me with intellectual sustenance throughout the lockdown years. Many of the ideas in this thesis can be traced back to the hours of Zoom meetings that we had over the course of that project. A note of thanks also to Mark Meenan who kept his cool whenever I crashed the HPC.

From the wider PhD community at Glasgow, I would like to thank all those who have taken the same path and made what can otherwise be a lonely journey, a worthwhile and enjoyable experience. In particular Arthur Galichère, Ashraful H. Mahfuze, Daminao Truchet, Prateek C. Bhan, Andres A. Gavaldon, Rohan Chowdhury, Maksym Solodarenko, Kulnicha Meechaiyo, Timo Hummel and many, many more. Particularly, my fellow knights Johanna Tiedmann, Andrea Benecchi and last but not least Spyridon Lazarakis who over these last years has been the greatest office mate, mentor and friend.

I am also grateful to all the other economists whom I have met in countless seminars and conferences all across the world, who were always friendly, collegiate and fun to be around. Their willingness to listen, offer advice and discuss ideas has had a great impact on my development as a researcher and economist.

Finally, I would like to thank my families by chance and by choice. My parents, Ralf & Bettina for raising me and encouraging me to pursue my interests. My siblings Sam & Greta for always providing the greatest company, particularly at Christmas time. Liz, for always being there for me throughout all these years, providing unconditional love, support and encouragement. And Frank, whose patience and silent wisdom is an example to us all.

INTRODUCTION

Motivation

Material inequality is a central feature of all human societies throughout history and - at least as far as we can know - pre-history (Scheidel (2017), Graeber & Wengrow (2021)). The pervasiveness of economic inequality and the differentials of power, prestige and influence that accompany it have been of interest to early political thinkers (e.g. Rousseau (1754), de Tocqueville (1835)) and economists (e.g. Smith (1759), (2010), Bentham (1843), and Mill (2022)).

Whilst inequality is a pervasive feature of all developed societies, attitudes towards differences along racial, gender and socioeconomic lines have become increasingly pro-egalitarian, as a result of a prolonged process of cultural and ideological development (see Holland (2019) for a discussion of the origins of these trends).

An important factor shaping societal attitudes with respect to economic inequality is the perceived fairness of the origins of differences in economic outcomes (Rawls (2020)). We often think that inequality is justified when it is based on differences in skill, talent, or effort, while we are less likely to accept differences that result from luck, the exploitation of unfair advantage or privilege.

This thesis focuses on a quantitative analysis of important determinants of labour income and wealth inequality: namely, the importance of human capital or skills, their accumulation and distribution; the role of technological change in determining the demand and supply of multidimensional skills; the contribution of wealth to intergenerational human capital accumulation and resilience in the face of large economic shocks; and ultimately how skills and luck create income and wealth inequality in a stochastic environment.

I focus on labour income and wealth inequality and how both interact with heterogeneity in human capital or skills inequality. For most individuals and households labour income is the main source of income and therefore changes to its distribution have the most direct effect on living standards. Wealth inequality is quantitatively the largest and most salient form of economic inequality and the dimension of economic inequality that is most prominently depicted in the media. All of these dimensions of inequality are important drivers informing political decision-making and societal debates around issues of equality, fairness and opportunity.

Human capital inequality and its interactions with technology and other economic phenomena is an important driver of both labour income and wealth inequality. Yet, many of the channels through which human capital affects economic inequality have so far not been analysed quantitatively. Some of these channels are described in more detail below and in the chapters of this thesis.

On a more philosophical note, human capital plays an important role in how we perceive fairness in relation to economic inequality. We generally believe that those who work harder, or smarter should be rewarded for the contributions they make to the larger society. However, in practice, it is almost always impossible to clearly separate these meritocratic elements from those that are generally deemed outside of an individual's control, be it family wealth and status, differences in opportunities or simple luck. Particularly wealth and human capital inequality are inextricably linked, feeding into and off each other in an eternal cycle of cause and effect. For example, individuals from wealthy families might be able to access better education and therefore accumulate more human capital than individuals from families with lower socioeconomic status, resulting in higher incomes and perpetuating differences across generations.

In this thesis, I focus my analysis on the United Kingdom. While most studies of income and wealth inequality focus on the United States, the UK has high levels of income and wealth inequality and has experienced transformative economic changes over the last 2 generations (see for example Blundell & Etheridge (2010), Belfield et al. (2017) for income inequality and Hills et al. (2013) for wealth inequality). More recently the Covid-19 pandemic has highlighted the importance of understanding inequalities for the economic recovery post-Covid and resilience to future shocks (Blundell et al. (2022), Marmot et al. (2021)).

For my analysis, I use quantitative economic models that allow for heterogeneity and where inequality is both a contributing and an endogenous outcome of the decisions made by the agents. These models build on and extend a large literature in Economics. Specifically, this thesis relates to three distinct modelling frameworks: i) Heterogenous agent incomplete market models for the analysis of wealth inequality following the seminal contributions of Aiyagari (1994), Bewley (1986), Huggett (1993) & Imrohoroglu (1989); ii) Models of wage determination where multidimensional skills are matched with differentiated tasks (c.f. Autor et al. (2003), Autor & Handel (2013), Sanders & Taber (2015), Roys & Taber (2016)); iii) Occupational choice models, where individuals with heterogenous abilities make utility-maximizing decisions (Roy (1951), Train (2009)). The specific modelling frameworks and the contributions provided by this thesis are discussed in more detail below and the specific chapters.

In order to calibrate these quantitative models, I use data from detailed micro surveys and household panels to estimate the underlying dynamic processes and evaluate the models' fit with respect to the economic reality they are supposed to describe. Amongst the data sources I use are the UK Household Longitudinal Study ("Understanding Society", University of Essex (2022)), the UK Quarterly Labour Force Surveys (Office for National Statistics (2023)) and the Skills and Employment Surveys (Felstead et al. (2019)).

The remainder of this introduction covers a review of the three chapters, as well as the methods employed in this thesis.

Review of Chapters

Chapter 1 explores the relationship between technological change, income and wealth inequality using a general equilibrium incomplete markets model. Income and wealth inequality have become increasingly prominent economic and political issues in recent years, particularly in the UK and other developed economies. It is widely accepted that income and wealth are interlinked, with those with higher incomes being able to save and invest more, potentially leading to reduced intergenerational mobility and reduced equality of opportunity.

For the most part, economists have approached these two dimensions of inequality separately: A large literature attributes changes in the distribution of incomes to technological change (c.f. Acemoglu & Autor (2011), Goldin & Katz (2007), Katz & Murphy (1992), Tinbergen (1974, 1975)). Technological innovations alter the relative demand for different types of skills and thereby change the distribution of income in the economy.

On the other point, the canonical model of wealth inequality - the incomplete markets model following the seminal contributions of Aiyagari (1994), Bewley (1986), Huggett (1993) & Imrohoroglu (1989), suggests that the degree of income inequality (or risk) is a key determinant of the level of precautionary savings and therefore wealth inequality. Taken together it seems only logical that a link between technological change and wealth inequality should be mediated by the former's effect on income inequality, but until recently this connection has rarely been explicitly addressed or quantitatively explored (see Moll et al. (2022) for a recent contribution to this debate).

In this chapter I develop a macroeconomic model that combines insights from both strands of the literature, to explore the quantitative importance of the rapid adoption of Information and Communication Technologies on the joint rise of income and wealth inequality in the UK since the 1980s. To this end I replace the standard reduced form earnings process commonly employed in heterogeneous agent incomplete market models with a process based on the "Skill-Weights-Approach" (Laezar (2009)) that has found application in the literature on Task Biased Technological Change (TBTC) (c.f. Autor et al. (2003), Autor & Handel (2011), Sanders & Taber (2012)). In the model, income depends on the skill set and occupation match of a worker, while the returns to skills depend on the supply-demand relationship in general equilibrium. This allows me to introduce a direct link between exogenous technological forces and income inequality within the model, to assess the endogenous response of wealth inequality to an increase in the demand for cognitive skills.

I calibrate a stationary version of the model, using a sample of workers from the UK's Understanding Society longitudinal household dataset, estimating the returns to cognitive and physical skills for different occupations, and calibrating a number of exogenous processes to match key parameters of the autoregressive earnings process for the UK. I then introduce the notion of cognitive skill-biased technological change, using differential rates of ICT adoption across different occupational groups to proxy the growing demand for cognitive skills over the period 1980-2016.

The model results suggest that this exogenous technological change can account for around 90% of mean wage growth and 80% of the growth of income inequality, and generate changes in wealth inequality that are consistent with those observed in the UK over the time period. Overall, the chapter provides evidence that cognitive skill-biased technological change has been a strong driver of the growth in income inequality and wage growth since the 1980s and that the presented hybrid model is a viable approach towards explaining the general trend of increasing income & wealth inequality.

Chapter 2 dives deeper into the relationship between wealth and human capital. I develop a novel heterogeneous agent incomplete markets model that features endogenous savings and risky human capital which workers accumulate during their working lives (see also Krebs (2002)). Wealth plays a key role in influencing the distribution of human capital and earnings, and, in addition to ability and opportunity, can have a huge impact on an individual's economic prospects. Recent research has emphasized the role of family wealth in the educational outcomes of young adults (see Lovenheim (2010), Karagiannaki (2017), Dräger (2022)), but rapid technological changes imply that skill accumulation stays important even after labour market entry. The need to constantly adapt to new technologies and work environments is putting a bigger emphasis on "post-formal" education and training, opening up a channel for wealth to play a role in affecting the outcomes of even prime-aged workers.

In general, richer workers can afford to invest more resources into human capital accumulation (c.f. Ben-Porath (1967)), and are therefore better able to take advantage of upcoming opportunities that might require specialised skills, and are also better able to maintain their current skillset in response to adverse shocks. Poorer individuals on the other hand may not have the available leeway to make such investments even if they would net a positive return in the long run. The absence of formal credit markets for education beyond the tertiary education sector likely exacerbates these issues since poor and borrowing-constrained individuals find it impossible to finance human capital investments with large up-front costs. Jointly, these dynamics imply that wealth becomes an important driver of human capital inequality, increasing earnings disparities over time and leading to a vicious cycle of self-reinforcing inequality. Correspondingly, a better understanding of the complex interactions between wealth, income and human capital can help inform policies that are tailored to the needs of different groups, reduce wealth disparities and ensure that everyone has access to the resources and opportunities needed to build and maintain their human capital.

I calibrate the model to match key features of the distribution of earnings, wealth and human capital in the UK prior to the Covid-19 pandemic. A quantitative analysis suggests that there exists a trade-off between wealth and human capital in terms of the level of investment in skills and knowledge and saving using a riskless asset. Workers with low levels of wealth invest considerably less in accumulating human capital than their counterparts with more initial savings. This phenomenon is particularly apparent at low levels of human capital when marginal returns to human capital investments are high. This has serious implications for economic development and social mobility.

I then analyse the response of the model to a uniform unexpected productivity shock, in order to investigate how households with different. The analysis shows that the initial distribution of wealth and human capital matters considerably for the aggregate response to such an aggregate shock, with households close to the borrowing constraint exhibiting the strongest negative reaction. This is because such households have fewer resources available to them to cushion the economic impact of the shock. This can lead to persistent increases in human capital and earnings inequality, with consequences for future economic growth and development.

Finally, I analyse the medium-run impact of the Covid-19 pandemic in the UK, demonstrating the need for public insurance policies to protect the most vulnerable in society. I find that without additional support measures, low-wealth households would have been most affected by the negative implications of the Covid-19 pandemic, leading to higher levels of human capital and wealth inequality for decades to come. Successful policy intervention has likely prevented welfare losses on the order of around 1.15% of lifetime consumption for the average household. An alternative UBI-type support system is shown to be slightly more welfare efficient but does not command a democratic majority amongst households. This highlights the importance of public insurance policies to protect the most vulnerable in society during large economic shocks like the ones the UK has been experiencing in recent years.

Chapter 3 investigates the changing skill supply distribution of university graduates in the UK. For most individuals, a university degree is the highest educational qualification they will receive, signalling that the holder possesses certain skills and abilities that will (hopefully) make them highly productive when they enter the labour force. Over the course of the 20th and early 21st century, the demand for a university education has been rising, showcased both by the stark increases in enrollment and graduation rates as well as the high premium that graduates can command for their labour once they join the labour force (c.f. Goldin & Katz (2007), Katz & Murphy (1992)).

In recent years over 2 million young people were enrolled in higher education institutions in the UK alone. But unlike a few generations ago, a degree no longer guarantees a successful professional career. Rising graduate earnings inequality and underemployment amongst university graduates (c.f. Altonji et al. (2016), Holmes & Mayhew (2016), Lindley & MacIntosh (2015)) means that for a rising proportion of young people, a degree looks more like an expensive gamble than a sound investment.

These observations have prompted much discussion in policy circles interested in ensuring that publicly funded higher education provides value for money and given rise to a large economic literature that investigates potential causes of these differences in outcomes (e.g. Altonji et al. (2016), Andrews et al. (2022) and Lovenheim & Smith (2022)). Much of this literature focuses on the role of subject of study ("college major") and the role that it plays in determining a graduate's career prospects. Subject of study is highly relevant for the types of skills a student will accumulate throughout their university education. An economics graduate will possess a vastly different set of skills from a freshly minted junior doctor, leading both to choose different occupations and setting them up for radically different career trajectories.

Yet despite its importance, the qualitative dimension of the skills supplied by university graduates has not been widely studied until now - in large part due to the problem that skills are mostly unobservable. I aim to overcome this challenge, by specifying and estimating an economic model of occupation choice for recent university graduates in the UK. Graduates draw their endowment of mathematical and verbal skills from subject-time-specific distribution functions upon graduation. Occupations in the labour market operate with different technologies, and therefore offer different wages to graduates with different levels of mathematical and verbal skills, resulting in an assignment of graduates to occupations based on the supply of and demand for specific multidimensional skills. By modelling this mechanism I am able to estimate the parameters that characterise the subject-specific multidimensional distribution functions that characterise the graduate skill supply and track their evolution over time.

Using data from the UK's Labour Force Survey from 1994 to 2019 combined with information on occupational skill requirements from the UK Skills and Employment Surveys, I structurally estimate the model to recover changes in the multidimensional skill distribution over a period of rapid change both in terms of the skills demanded in the labour market, as well as the structure of the UK education sector as a whole. I find that on average, the median graduate's endowment of mathematical skills has increased, while verbal skills have decreased. Over the same period, skill inequality as measured by the Gini coefficient for mathematical skills has decreased, while the Gini coefficient for verbal skills has increased slightly.

These large changes in the skill supply not only have important implications for individuals who are entering university but also for the wider economy as a whole. Universities operate at the frontiers of knowledge creation and it is crucial that graduates are able to absorb relevant insights and transport these skills into the labour market where they can aid the adoption and development of new technologies and ultimately support a growing and prosperous economy. As technological change accelerates the actual types of skills that individuals have and how these interact with a rapidly changing technological environment become increasingly relevant. This chapter provides a theoretical and empirical framework for thinking about multidimensional skill heterogeneity and how the qualitative dimension of the skill supply can change in response to both institutional reforms and changing demand conditions.

Methodology

All chapters presented in this thesis present cases of quantitative economic models that are calibrated or estimated using microdata to investigate different aspects of wealth and human capital inequality. Human capital is often difficult to measure making the use of theoretical frameworks essential to the study of this important variable. Chapters 1 & 2 present models that are variations on the standard heterogeneous agent incomplete markets framework following the work of Aiyagari (1994), Bewley (1986), Huggett (1993), Imrohoroglu (1989). Chapter 3 presents a custom occupational choice model, that combines aspects from both the wider returns to tasks literature (c.f. Autor et al. (2003), Autor (2013), Sanders & Taber (2012)) and consumer choice modelling (c.f. Train (2009)).

Chapter 1 investigates the role that cognitive skill-biased technological change has played in explaining the joint increase in income and wealth inequality experienced in the UK since the 1980s. The Aiyagari (1994) model is the standard workhorse model for analysing wealth inequality with numerous applications in macroeconomic research (see Heathcote et al. (2009), or Quadrini & Rios-Rull (2015)). The key aspects of this class of model are: i) a continuum of agents instead of a representative household, allowing the model to represent distributions of income and wealth; ii) a risky flow of resources to the agent for example in the form of uncertain labour income; iii) limited markets in the sense that the agent is unable to perfectly insure themselves against income risk and therefore needs to self-insure using precautionary savings.

The chapter's key contribution is to characterise the sources of income risk in the Aiyagari framework, and linking the earnings process to recognizable skills, thereby allowing an analysis of how the growing demand for cognitive skills has impacted inequality. The standard approach to modelling earnings risk would be to define an exogenous Markov process that provides an approximation to a continuous AR(1) income process that has been estimated on relevant cross-sectional or panel data. This process provides a reduced form view of the variation of incomes across the population in so far as that it combines all possible sources of income risk, including, but not limited to, differences in skills, changes in jobs, temporary and permanent health episodes, or simply luck. While such a process can be used to examine changes in the overall level of income risk, it is unable to provide a decomposition of the factors driving these changes.

In order to allow for such a decomposition, I propose a new exogenous income process, that combines different skills and differences in returns to these skills. Following the literature on tasks and skills (c.f. Autor et al. (2003), Autor (2013), Sanders & Taber (2012)), I assume that agents are endowed with two distinct types of skills: cognitive and physical. These skills are rewarded proportionately to how intensely they are used in the agent's current occupation. This means that the income process of the model is constituted of two interacting sub-processes: i) an exogenous process assigning agents to occupations; ii) an exogenous process determining the skill set of the agent.

I estimate the occupation-specific returns to cognitive and physical skills using Understanding Society data using available proxies for both physical and cognitive skills. I then estimate the parameters determining the exogenous skill processes and occupation assignment process using a combination of panel data and indirect inference ensuring that my proposed income process matches a standard AR(1) process.

Having incorporated this new process into an otherwise standard heterogeneous agent model, I can study how changes to the return of different skills affect income risk and in turn wealth inequality The main application of the model is to the increasing return to cognitive skills following the rapid adoption of Information and Communication Technologies (ICT) during the 1980' and 90's. I estimate occupation-specific computer adoption curves using survey data, to proxy the changes in the returns to cognitive skills. Then I simulate the transition of the model economy along the deterministic transition path capturing both changes in income risk and the distribution of wealth, brought about by this technological change.

Chapter 2 provides a more ambitious extension of the incomplete markets framework, by proposing a model in which agents can accumulate both wealth and human capital. This innovation relaxes the general assumption that workers have no control over the stream of future incomes and allows me to study the endogenous joint distribution of wealth and human capital. Despite the clear importance of human capital for income and income risk, endogenous human capital has rarely been studied in the context of wealth inequality. Important exceptions from this are for example Krebs (2003) who studies comparative statics in a stationary economy with risky human capital, and Huggett et al. (2011) who study life-cycle income profiles using a life-cycle model with wealth and human capital.

The important contribution of the model proposed in this chapter is to distinguish between an agent's stock of human capital and the productivity of that stock which depends on factors specific to the agent's employment. Prior modelling approaches are unable to make this distinction since they infer human capital residually, using measures of wages or incomes. Hence, these approaches to modelling risky human capital conflate changes to the stock of human capital, brought about by for example skill accumulation and changes in the return of these skills, due to for example a change in occupation or industry. Modelling these factors separately does not only bring the model closer to the way that economists generally think about human capital, but also allows me to investigate the effect of factors that affect only human capital accumulation from those that affect the productivity of skills.

The model assumes that agents face a risky exogenous process that determines the productivity of their current stock of human capital. I calibrate this process separately from the agent's human capital accumulation by matching both the observed earnings distribution and a proxy for the human capital distribution available in the UK's Understanding Society data, using a minimum distance procedure.

Agents receive labour income according to the productivity of their human capital stock. They can consume, save or invest in accumulating more human capital. This mechanism creates an important feedback mechanism between the distribution of wealth and the distribution of human capital and hence labour income in the economy. I show that in the stationary economy, even small differences in initial wealth can lead to large and persistent differences in human capital, amplifying income inequality. This mechanism shows that there is an important feedback from wealth distribution to earnings inequality that has not been addressed by standard models.

I also show that agents with low levels of wealth are more exposed to aggregate shocks. I implement a surprise TFP shock in the model economy and find that households close to the borrowing constraint, see a dramatic reduction in their human capital accumulation as a result of the shock, leading them to have persistently lower levels of human capital, increasing income inequality and slowing down the economy's recovery. This result runs somewhat counter to the famous result by Krussell & Smith (1998), that the distribution of wealth does not matter significantly for the response of the economy to aggregate shocks, and suggests that potentially human capital is an important channel in facilitating fluctuations in models with aggregate risk.

For the main application of the model, I turn to the recent Covid-19 pandemic. I calibrate a shock to the model economy to reproduce the productivity losses experienced during the period of lockdowns and implement a limit on consumer spending in line with the experience throughout the pandemic years. I simulate the dynamic transition of the economy in response to the Covid-19 shock and evaluate both the medium-run consequences and the effectiveness of different policy interventions.

Chapter 3 presents an empirical investigation into the changing skill content of different university degrees. A university degree constitutes one of the highest educational qualifications an individual can attain, signalling that the degree holder has obtained a variety of specialised skills. While skills, their nature and their distribution amongst the population are questions of great relevance to a society little is known about the types of skills that university graduates possess, and how these might have changed in response to changing institutional setups and changes in the demand for different types of skills.

This lack of evidence is primarily driven by the particular challenges that skills pose to empirical analysis: skills are largely unobservable and tend to be correlated with other labour market characteristics, making the estimation of skills or their prices difficult (see Roy (1951) who lays out the problem rather eloquently). The chapter's main contribution is to exploit the selection effect, in order to use observable labour market characteristics to infer unobservable skills.

To this end, I develop a novel structural economic model of occupational choice for university graduates. Following the literature on multidimensional skills and tasks (c.f. Autor et al. (2003), Autor (2013), Sanders & Taber (2012)), I assume that different occupations have different skill requirements and hence a worker's wage in different occupations will depend on how well their skill set matches the specific requirements of the occupation. If new entrants into the labour market can observe their own skill set and also have knowledge of the skill requirements of different occupations, their optimization behaviour will lead them on balance to select jobs that make the best use of their skills. So, far from being a problem for an estimation, selection actually plays a crucial role here since it provides an indication to the nature of the skills the worker has.

In the chapter, I discuss how the model can then be estimated using simulated maximum likelihood, and provide an algorithm for performing the estimation. For the application, I focus on quantifying changes to the distribution of skills amongst recent university graduates in the UK over the last 25 years. Focusing on two relevant types of skills - mathematical/technical and verbal/organisatorial - the model provides estimates of subject-specific skill distributions allowing a quantitative assessment of the skills of a typical graduate as well as the degree of skill inequality between graduates.

For the estimation, I use data on university graduates from the UK Labour Force Survey, as well as proxies for the skill requirements of different occupations derived from the UK Skills and Employment Surveys.

Taking a long view, the chapter provides separate estimates for three time periods, covering the period from 1994 until 2019. The paper finds that between 1994 and 2019, the typical graduate's level of mathematical skills increased by 140% while verbal skills decreased by close to a third. Looking closer at 5 different major subject categories, I find that this trend is driven by increasing specialization for STEM and Business & Economics degrees and increasing generalization among Arts & Humanities and Other Subjects. For most graduates mathematical/technical skills have become the single biggest contributing factor to their earnings, making up around 50% of their hourly wage compared to 27% in the mid-90s. Counterfactual simulations suggest that in the absence of changes to the subject-specific skill distributions, mean wages would be up to 8% lower, while wage inequality would be up to 5% larger. The results suggest that graduate skill supply has adjusted to changing labour market requirements.

BIBLIOGRAPHY

[1] Acemoglu, D. and Autor, D. "Skills, tasks and technologies: Implications for employment and earnings," Handbook of labor economics (4), 2011, pp. 1043–1171.

[2] Altonji, Joseph G., Peter Arcidiacono, and Arnaud Maurel. "The analysis of field choice in college and graduate school: Determinants and wage effects." In Handbook of the Economics of Education, vol. 5, pp. 305-396. Elsevier, 2016.

[3] Aiyagari, S. Rao. "Uninsured idiosyncratic risk and aggregate saving." The Quarterly Journal of Economics 109, no. 3 (1994): 659-684.

[4] Andrews, Rodney J., Scott A. Imberman, Michael F. Lovenheim, and Kevin M. Stange. The Returns to College Major Choice: Average and Distributional Effects, Career Trajectories, and Earnings Variability. No. w30331. National Bureau of Economic Research, 2022.

[5] Autor, David, H. "The "task approach" to labor markets: an overview." Journal for Labour Market Research 46, no. 3 (2013): 185-199.

[6] Autor, D. H., Levy, F. and Murnane, R. J. "The Skill Content of Recent Technological Change: An Empirical Exploration," The Quarterly Journal of Economics (118:4), 2003, pp. 1279–1333.

[7] Autor, D. H. and Handel, M. J. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages," Journal of Labor Economics (31:S1), 2013, pp. S59–S96.

[8] Belfield, Chris, Richard Blundell, Jonathan Cribb, Andrew Hood, and Robert Joyce. "Two decades of income inequality in Britain: the role of wages, household earnings and redistribution." Economica 84, no. 334 (2017): 157-179.

[9] Ben-Porath, Yoram. "The production of human capital and the life cycle of earnings." Journal of political economy 75, no. 4, Part 1 (1967): 352-365.

[10] Bentham, Jeremy. The Works of Jeremy Bentham. Vol. 7. W. Tait, 1843.

[11] Bewley, Truman. "Stationary monetary equilibrium with a continuum of independently fluctuating consumers." Contributions to mathematical economics in honor of Gérard Debreu 79 (1986).

[12] Blundell, Richard, and Ben Etheridge. "Consumption, income and earnings inequality in Britain." Review of Economic Dynamics 13, no. 1 (2010): 76-102.

[13] Blundell, Richard, Monica Costa Dias, Jonathan Cribb, Robert Joyce, Tom Waters, Thomas Wernham, and Xiaowei Xu. "Inequality and the COVID-19 Crisis in the United Kingdom." Annual Review of Economics 14 (2022): 607-636.

[14] De Tocqueville, Alexis. "Democracy in America [1835]." (1956).

[15] Dräger, Jascha. "The role of parental wealth in children's educational pathways in Germany." European Sociological Review 38, no. 1 (2022): 18-36.

[16] Felstead, A., Gallie, D., Green, F., Henseke, G. Skills and Employment Surveys Series Dataset, 1986, 1992, 1997, 2001, 2006, 2012 and 2017. [data collection]. UK Data Service. SN: 8589, DOI: 10.5255/UKDA-SN-8589-1 (2019).

[17] Goldin, C. and Katz, L. F. "The race between education and technology: the evolution of US educational wage differentials, 1890 to 2005", Technical report, National Bureau of Economic Research, National Bureau of Economic Research, 2007

[18] Graeber, David, and David Wengrow. The dawn of everything: A new history of humanity. Penguin UK, 2021.

[19] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. "Quantitative macroeconomics with heterogeneous households." Annu. Rev. Econ. 1, no. 1 (2009): 319-354.

[20] Hills, John, Francesca Bastagli, Frank Cowell, Howard Glennerster, and Eleni Karagiannaki. Wealth in the UK: distribution, accumulation, and policy. Oxford University Press, 2013.

[21] Holland, Tom. Dominion: The making of the Western mind. Hachette UK, 2019.

[22] Holmes, Craig, and Ken Mayhew. "The economics of higher education." Oxford review of economic policy 32, no. 4 (2016): 475-496.

[23] Huggett, Mark. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." Journal of economic Dynamics and Control 17, no. 5-6 (1993): 953-969.

[24] Huggett, Mark, Gustavo Ventura, and Amir Yaron. "Sources of lifetime inequality." American Economic Review 101, no. 7 (2011): 2923-2954.

[25] Imrohoroğlu, A. "Cost of business cycles with indivisibilities and liquidity constraints," Journal of Political economy (97:6), 1989, pp. 1364–1383.

[26] Karagiannaki, Eleni. "The effect of parental wealth on children's outcomes in early adulthood." The Journal of Economic Inequality 15 (2017): 217-243.

[27] Katz, L. F. and Murphy, K. M. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors", 1992, pp. 35-78.

[28] Krebs, Tom. "Human capital risk and economic growth." The Quarterly Journal of Economics 118, no. 2 (2003): 709-744.

[29] Krusell, Per, and Anthony A. Smith, Jr. "Income and wealth heterogeneity in the macroeconomy." Journal of political Economy 106, no. 5 (1998): 867-896.

[30] Lazear, E. P. "Firm-Specific Human Capital: A Skill-Weights Approach," Journal of Political Economy (117:5), 2009, pp. 914–940.

[31] Lindley, Joanne, and Steven McIntosh. "Growth in within graduate wage inequality: The role of subjects, cognitive skill dispersion and occupational concentration." Labour Economics 37 (2015): 101-111.

[32] Lovenheim, Michael F. "The effect of liquid housing wealth on college enrollment." Journal of Labor Economics 29, no. 4 (2011): 741-771.

[33] Lovenheim, Michael F., and Jonathan Smith. Returns to different postsecondary investments: Institution type, academic programs, and credentials. No. w29933. National Bureau of Economic Research, 2022.

[34] Marmot, Michael, Jessica Allen, Peter Goldblatt, Eleanor Herd, and Joana Morrison. "Build Back Fairer." London: Institute of Health Equity, Available online: https://www. health. org. uk/sites/default/files/upload/publications/2020/Build-back-fairer-the-CO VID-19-Marmot-review. pdf (accessed on 16 March 2022) (2021).

[35] Mill, John Stuart. The Collected Works of John Stuart Mill. DigiCat, 2022.

[36] Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo. "Uneven growth: automation's impact on income and wealth inequality." Econometrica 90, no. 6 (2022): 2645-2683.

[37] Office for National Statistics. Labour Force Survey. [data series]. 7th Release. UK Data Service. SN: 2000026, DOI: 10.5255/UKDA-Series-2000026 (2023).

[38] Rawls, John. A theory of justice: Revised edition. Harvard university press, 2020.

[39] Rousseau, Jean-Jacques. "Discourse on the origin and basis of inequality among men." Rousseau J.-J. Traktaty (1754): 31-109.

[40] Roy, Andrew Donald. "Some thoughts on the distribution of earnings." Oxford economic papers 3, no. 2 (1951): 135-146.

[41] Roys, Nicolas, and Christopher Taber. "Skills prices, occupations and changes in the wage structure." (2016).

[42] Sanders, C. and Taber, C. "Life-Cycle Wage Growth and Heterogeneous Human Capital," Annu. Rev. Econ. (4:1), 2012, pp. 399–425.

[43] Scheidel, Walter. "The great leveler." In The Great Leveler. Princeton University Press, 2017.

[44] Smith, Adam. "An inquiry into the nature and causes of the wealth of nations: Volume One." London: printed for W. Strahan; and T. Cadell, 1776., 1776.

[45] Smith, Adam. The theory of moral sentiments. Penguin, 2010.

[46] Tinbergen, J. "Substitution of graduate by other labour," Kyklos (27:2), 1974, pp. 217–226.

[47] Tinbergen, J. "Substitution of academically trained by other manpower," Review of World Economics (111:3), 1975, pp. 466–476.

[48] Train, Kenneth E. Discrete choice methods with simulation. Cambridge university press, 2009.

[49] University of Essex, Institute for Social and Economic Research. (2022). Understanding Society: Waves 1-12, 2009-2021 and Harmonised BHPS: Waves 1-18, 1991-2009. [data collection]. 17th Edition. UK Data Service. SN: 6614, http://doi.org/10.5255/UKDA-SN-6614-18.

[50] Quadrini, Vincenzo, and José-Víctor Ríos-Rull. "Inequality in macroeconomics." In Handbook of income distribution, vol. 2, pp. 1229-1302. Elsevier, 2015.

CHAPTER 1

COGNITIVE SKILL-BIASED TECHNOLOGICAL CHANGE, INCOME & WEALTH INEQUALITY IN THE UK

1 Introduction

In recent years, economists and laypeople alike have become increasingly focussed on inequality as a defining economic and political issue.[1] In Britain, for example, income and wealth inequality seem to have increased for the last decades (c.f. Hills et al. (2013), Roine & Waldenstroem (2015), Belfield et al. (2017)), leading to increasing concerns about the implications for intergenerational mobility, equality of opportunity and distributional justice.

Economists working on the evolution of income inequality have largely focussed on the importance of technological change and the changing supply of skills,[2] most aptly summarized by Jan Tinbergens (1974, 1975) *"Race between Education and Technology"* (c.f. Acemoglu & Autor (2011), Katz & Murphy (1992), Goldin & Katz (2007)). Their insights suggest, that incomes are affected by a complex interplay of the technology-driven demand for and the supply of differentiated skills. Of particular relevance appears

[1] See for example the huge range of popular books that deal with the subject matter, such as Atkinson (2015), Milanovic (2016), Piketty (2013), Stiglitz (2012), and many more.

[2] There are of course other prominent explanatory frameworks, such as the majorily trade based Heckscher-Ohlin model, or models based on political economy channels (see Vivarelli (2014) for an overview).

to be the rapid adoption of advanced Information and Communication Technologies (ICT) throughout the 1980s and 90s that led to increases in the demand for high-level abstract reasoning skills, and replaced many routine work tasks across many developed economies.

Changes to the wage structure of the economy have important implications for wealth accumulation and distribution, by changing the nature and scale of income risk faced by workers and households. Most recent work on wealth inequality, emphasizes the role that precautionary savings play in generating differences in wealth across households, an insight that is condensed in the incomplete markets models of the Bewley-Huggett-Aiyagari variety (c.f. Aiyagari (1994), Bewley (1986), Huggett (1993), Imrohoroglu (1989)). In these models uncertainty about future income induces individuals to accumulate precautionary savings, leading to an endogenous distribution of wealth. Over the last decades, these models have become the benchmark for studying wealth inequality and found numerous applications to different macroeconomic problems (see Heathcote et al. (2009), or Quadrini & Rios-Rull (2015) for a survey).

The connection between these two strands of literature should be of interest to economists studying either phenomenon: If technology generates income inequality, and income inequality (risk) is a main component in models of wealth inequality, then the case can be made that both phenomena are connected.

The aim of this chapter is to build a model that captures the spirit of this idea and investigate the question of to what extent the rise in income and wealth inequality can be attributed to the advent of ICT in the 1980s. To this end, I insert a technology-driven model of income determination, that is based on the *"Skill-Weights-Approach"* (Laezar (2009)) into a standard Aiyagari (1994) model.

In this case, income depends on the skill set and occupation match of a worker while the returns to skills depend on the supply-demand relationship in general equilibrium. For simplicity, the constituent parts of the income process remain purely exogenous, as in a standard HIM application, but by taking a step back and modelling the relationship between the demand and supply of different skills explicitly, the model can shed some light on the potential effects of differentiated technological change. The aim of this is not to create a model that perfectly captures the determinants of income but to add just enough complexity so that the model can capture the theoretical entities, and still be simple enough to be calibrated with available data and solved easily with known techniques.

After calibrating a stationary version of the model, using a sample of workers from the UK's *Understanding Society* dataset, I use the model to explore the response of the endogenous variables to heterogeneous technological change, as computers penetrate different occupational niches at varying rates. In this, the model is uniquely suited for the study of income and wealth inequality dynamics over the short to medium term, as it allows me to trace the effects of technological change on both variables simultaneously.

The results suggest that the increasing adoption of computers and computerized equipment over the period $1980 - 2016$ can account for around 90% of mean wage growth and 80% of the growth of income inequality. Furthermore, the model's predictions are in line with stylized facts about the evolution of wealth inequality. This suggests, that cognitive skill-biased technological change might be a good approach towards explaining the general trend of increasing income & wealth inequality for a large sample of the working population in the UK.

1.1 Related Literature

The chapter relates to roughly four different strands of literature:

Firstly, it relates to the literature on technological change in heterogenous agent models: In a recent paper Moll et al. (2022) study the effect of technological change on wealth inequality in the US. Their proposed mechanism is that technological change raises the return to wealth and thereby amplifies wealth inequality. This chapter looks at a different channel, namely changes to the returns to different skills, and how these might result in higher levels of wealth inequality. Somewhat related is Hubmer et al. (2017), who study potential drivers of changes in wealth inequality in the US, but do not explicitly consider changes in technology.

Secondly, the chapter relates to the literature on Task Biased Technological Change (TBTC) (c.f. Acemoglu & Autor (2012), Acemoglu & Restrepo (2015)), in so far that it embraces the notion that technological change can have different effects depending on an individual's comparative advantage. Workers are subject to two exogenous forces that determine their idiosyncratic labour productivities: their skill set and their occupation match. As the technological environment changes due to *cognitive skill-biased technological change,* different skill-occupation states become more productive whilst others fail to keep up, leading to an endogenous increase in wage inequality. An important difference to these frameworks is that my model abstracts from the endogenous assignment of workers to tasks and instead focuses on the feedback effect between income inequality, precautionary savings behaviour and capital accumulation. The chapter also contributes to the study of the effects of computerization, by showing that the increasing return to cognitive skills associated with increased computer usage is sufficient to explain the bulk of the growth in income inequality and wage growth since the 1980s, as well as some stylized facts about growing wealth inequality.

Thirdly, the chapter contributes to the literature on between- and within-group inequality and labour market polarization (c.f. Angelopoulos et al. (2017), Cortes (2016), Goos & Manning (2003), Goos et al. (2014), Kambourov & Manovski (2009)), by offering a simple and attractive framework for studying these topics under different technological change scenarios. In the same vein, the chapter extends research into the relationship between changes to income uncertainty and skill premia (c.f. Heathcote et al. (2010), Guvenen & Kuruscu (2012), Slavik & Yazici (2018)), by generating meaningful between-group wage premia from a simple technological process. As the model is rooted in the task-skill literature, the resulting group distinctions are occupation based, rather than the more familiar distinctions based on education, even though some similarities might be inferred. Similarly, wage premia refer to occupation wage premia rather than *skill* or *degree premia.* Extending the model to include these additional distinctions might be an interesting extension.

Fourthly, the chapter contributes to the literature on Tasks and Skills (c.f. Autor et al. (2003), Autor (2013), Sanders & Taber (2012)) and in particular the *"Skill-Weights"* literature (c.f. Autor & Handel (2013), Laezar (2009)), by providing an alternative way of calibrating skill weights using directly observed skill measures.

The *"Task-Skill"* framework (c.f. Autor et al. (2003), Autor (2013), Sanders & Taber (2012)), suggests that productivity depends crucially on the match between a worker's multidimensional skill set and their job's multidimensional skill demand. A lot of recent work in this area (c.f. Yamaguchi (2012), Guvenen et al. (2017), Lise & Postal-Viney (2017)) has already shown that this approach can be successfully employed to explain life-cycle income profiles.

Traditionally, skill weights are calibrated using task survey data, such as is provided by the US O*Net database. Here I use available measures of skills to infer the return to the corresponding tasks via econometric analysis. This approach avoids the selection problem, inherent in models of occupational choice, and provides some evidence for the validity of common survey-based approaches (c.f. Autor & Handel (2013), Bisello (2013), Firpo et al. (2011), Gathman & Schoenberg (2010), Rohrbach-Schmidt & Tiemann (2013)).

The rest of the chapter is structured as follows: Section 2 presents the production environment of the economy and elaborates on the role of differentiated skill supplies & demands; Section 3 presents the household's problem; Section 4 gives the definition of the recursive, stationary equilibrium; Section 5 provides a simple illustrative example; Section 6 describes the calibration of the model; Section 7 compares the properties of the model generated income process with the data and outlines the implications for wealth inequality; Section 8 provides a decomposition of income risk into cognitive and physical skill risk; Section 9 presents the dynamic effects of cognitive skill-biased technological change due to increased computerization & Section 10 concludes.

2 Production Environment

2.1 Representative Firm

There is a representative, competitive firm that produces a single homogenous good (Y) using a customary Cobb-Douglas production technology using capital (K), as well as labour services (L) as inputs:[3][4]

$$Y = K^{\alpha} L^{1-\alpha} \tag{1.1a}$$

$$\alpha \in (0, 1) \tag{1.1b}$$

Labour services L are differentiated and derived from the combination of intermediate labour services l_n, supplied by different occupations, or *occupation islands,* denoted by $n = \{1, .., N\}$:

$$L = \sum_{n=1}^{N} l_n \tag{1.2}$$

On these islands, intermediate labour services l_n are produced by a weighted linear combination of a number of general skills (e.g. Manual, Cognitive, Interpersonal, etc.), denoted by $m = \{1, .., M\}$.[5] S_m^n summarizes the total amount of skill m supplied by workers on the occupation island n.

$$l_n = \sum_{m=1}^{M} \lambda_m^n S_m^n \tag{1.3a}$$

$$\infty > \lambda_m^n \geq 0, \forall m, n \tag{1.3b}$$

The *occupation islands* can be thought of as different departments of the representative firm, that perform different work activities, each using a specific production technology, contributing to the production of the homogenous output. The firm has no influence over the allocation of workers across these departments but is able to purchase units of skills directly from each island and pay differentiated wages.

[3]Since all aggregate quantities are constant in the stationary equilibrium, time subscripts are omitted in this part of the exposition.

[4]There is no aggregate productivity term A in the production function since all relevant productivity is contained in the occupation-skill-specific productivity parameters.

[5]In a sense these skills correspond to the types of "Task Specific Human Capital" in Gathmann & Schonberg (2010).

Each occupation n is associated with a set of weights $\lambda^n = (\lambda_1^n, ..., \lambda_M^n)$.[6] The weights λ_m^n indicate how productive the general skill m is in the production process in occupation n. It has long been argued that different work tasks require different combinations of skills so that the same worker might be more or less productive, depending on which work activity he applies his skill set to (c.f. Autor et al. (2003), Acemoglu & Autor (2012)).

For example, an accomplished economist might find himself quite out of his depth, when he has to apply himself to work that requires less analytic and more manual skills.[7] The non-negativity constraint on the skill weights accounts for the fact that skills cannot have negative productivities - even though it is possible that some skills contribute very little, or even nothing to the productivity of an island.

Overall, this approach to labour productivity has sometimes been referred to as the "*Skill Weights Approach*" (see Laezar (2009)). The main benefit of this formulation is that it accounts for the fact that skills are sufficiently general to be transferred across occupations, whilst at the same time having different productivities in different occupations.[8]

Wages are paid per unit of skill and vary across the N occupations and M different skills, allowing for the possibility, that different skills are differentially rewarded in different occupations. In particular w_m^n is the amount paid for 1 unit of skill m from a worker currently employed in occupation n.

The competitive firm takes prices ($\{w_m^n\}_{\forall m}^{\forall n}$ and r) as given and chooses K which depreciates at rate δ, and the amount of each skill it wants to purchase from each island (S_m^n) to maximize profits Σ:

$$\max_{K,\{S_m^n\}_{\forall m}^{\forall n}} \Sigma = K^\alpha \left(\sum_{n=1}^{N} \sum_{m=1}^{M} \lambda_m^n S_m^n \right)^{1-\alpha} - (r+\delta)K - \sum_{n=1}^{N} \sum_{m=1}^{M} w_m^n S_m^n \tag{1.4}$$

Taking the derivative wrt. S_m^n gives:

$$\frac{\partial \Sigma}{\partial S_m^n} = (1-\alpha)\lambda_m^n K^\alpha \left(\sum_{n=1}^{N} \sum_{m=1}^{M} \lambda_m^n S_m^n \right)^{-\alpha} - w_m^n \tag{1.5}$$

[6] Sometimes I will refer to $\lambda = (\lambda^1, ..., \lambda^N)$ which denotes the set of all sets of weights.

[7] This all reaches back to the seminal work by Roy (1951).

[8] For example, the transferability of a skill k between two occupations x and y could be expressed by the Euclidian distance $|\lambda_k^x - \lambda_k^y|$. For a more sophisticated approach see e.g. Gathmann & Schoenberg (2010).

Noting that

$$\left(\sum_{n=1}^{N}\sum_{m=1}^{M}\lambda_m^n S_m^n\right) = \sum_{n=1}^{N} l_n = L \tag{1.6}$$

and setting the first order condition to 0 gives us the expression for the occupation-skill-specific wage:

$$w_m^n = \frac{(1-\alpha)\lambda_m^n K^\alpha}{L^\alpha} \tag{1.7a}$$

Which is the familiar wage equation for a Cobb-Douglas production function, scaled by the skill weight parameter λ_m^n.

Finally, the first-order condition for the demand for capital is the familiar expression:

$$r = \alpha\left(\frac{L}{K}\right)^{1-\alpha} - \delta \tag{1.8}$$

2.2 Skill Supply

Time is discrete and denoted by $t = 0, 1, 2...$. The economy is populated by a continuum of infinitely lived agents, distributed on the interval $I = [0, 1]$ with measure ζ and total mass 1. Agents are ex-ante homogenous and only differ with respect to the history of shocks they receive. All shocks are exogenous and take one of two forms:

1. Shocks to the skill endowment s_t that can be interpreted as events that affect a worker's capacity to work in both negative and positive ways. For example, a worker who suffered an accident could have his physical skills greatly reduced, while picking up a stimulating hobby might increase someone's mental abilities.

2. Shocks to the worker's occupation match n_t that are interpreted as occupation transitions.[9]

Both types of shocks affect a worker's productivity, albeit in differentiated ways.

In traditional treatments of income risk, shocks are typically not separated by their source, but rather by their persistence (c.f. Meghir & Pistaferri (2011)), if any distinction is made at all. In this chapter, the distinction is made deliberately and serves an important function. Technological change is thought to occur at the occupation level (c.f. Autor et al. (2003)) and thus "outside" of the worker. But if all that can be observed is the evolution of labour income based on the joint movement of skills and occupation matches it is impossible to identify the contribution of technological change to earnings risk. By separating two sources of income risk, and calibrating each independently, the model is

[9] The exogeneity of occupation transitions is clearly a strong assumption, however not any stronger than the all-to-common assumption of exogenous income processes. There are also many "exogenous" reasons why someone might change their occupation: necessitated by a change of location, flights of fancy...

able to identify the relevant effects and therefore give a more accurate account of the effect of technological change on income and wealth inequality.

There are a finite number of M different skills, each with an associated state space $\tilde{S}^m = [0,1]$[10]. Let us denote the Cartesian product of all these state spaces as: $\tilde{S} = \tilde{S}^1 \times ... \times \tilde{S}^M$ with associated $\sigma-$algebra $\tilde{\mathcal{S}}$.

Every period t, a worker $i \in I$ independently[11] draws a skill endowment $s_t = (s_{1,t}, ..., s_{M,t}) \in \tilde{S}$. In particular, I assume that these draws are described by a first-order Markov Process, with Transition Function $\Omega : \tilde{S} \times \tilde{\mathcal{S}} \Rightarrow [0,1]$. Throughout the chapter, I assume that this process has a *unique stationary distribution* $\bar{s}$ and impose that all draws are initiated from this stationary distribution.

The values of the worker's skill endowment s_t summarize his ability to perform different tasks, with values close to 1 indicating a high proficiency and values closer to 0 indicating low skills in that area. As s_t evolves over time, an individual worker's abilities change, however since $\bar{s}$ is a stationary distribution, the aggregate joint distribution remains unchanged.

The stationarity of $\bar{s}$ allows us to define the time invariant mean level of skill m across the population as:

$$\bar{\eta}_m = \int s_m \Omega(d\bar{s}) \tag{1.9}$$

Additionally, in every period individuals are assigned to one of the N *occupation islands*. For simplicity, the process of assignment is taken to be exogenous and follows a Markov Chain of order one, summarized by the occupation transition matrix Π. Assuming that Π is *irreducible* and *aperiodic*, each stochastic process (n, Π) induces a *unique stationary distribution* $\bar{\mu} = (\bar{\mu}_1, ..., \bar{\mu}_N)$ that describes the distribution of workers across the N occupations.[12] Again I assume throughout the chapter that we are always drawing from $\bar{\mu}$ whenever necessary.

Workers supply their skills inelastically on their respective islands. Since the processes for skills and occupation transitions are independent from one another, it follows, that the total supply of skill m to occupation n is given by:

$$S_m^n = \bar{\mu}_n \bar{\eta}_m \tag{1.10}$$

Which is the proportion of workers on island n multiplied by the mean skill level of skill m.

[10] This is a necessary, but arbitrary normalization. Any finite upper bound on the state space will suffice.

[11] Skill draws are independent across agents, but may be correlated across time.

[12] See for example Chapter 3.1 in Miao (2014).

This allows us to express equation (1.3a) as:

$$l_n = \sum_{m=1}^{M} \lambda_m^n \bar{\mu}_n \bar{\eta}_m \tag{1.11}$$

Noting that $\bar{\mu}_n$ is constant across all m, we can rewrite this as:

$$l_n = \bar{\mu}_n \sum_{m=1}^{M} \lambda_m^n \bar{\eta}_m \tag{1.12}$$

Finally, summing over all N occupation islands, we derive the expression for the aggregate measuring labour services:

$$L = \sum_{n=1}^{N} \bar{\mu}_n \sum_{m=1}^{M} \lambda_m^n \bar{\eta}_m \tag{1.13}$$

This suggests that for any given transition matrix Π and skill process $(\tilde{S}, \tilde{\not{S}}, \Omega)$, there exists a stationary distribution of skills over occupations that pins down the labour supply L and therefore the prices $\{w_m^n\}_{\forall m}^{\forall n}$ & r, for a given K. Furthermore, since the stationary distributions implied by Π and $(\tilde{S}, \tilde{\not{S}}, \Omega)$ are unique, then so are the derived quantities and prices.

The next section describes the choices of the representative household, acting in this environment.

3 The Household's Problem

Households discount the future at rate $\beta \in (0, 1)$ and inelastically supply their entire skill endowment to the *occupation island* they are assigned to. In return, they receive labour income ω which they can use to purchase the consumption good c or invest in a safe asset a that pays the return r. This asset corresponds to capital that is used in production (K) and the total capital stock of the economy is given by the sum of all savings held by the households. Households can also use their existing asset stock to finance further consumption or take on a limited amount of debt.

Financial markets in this economy are imperfect and individuals face a borrowing limit:

$$a_{t+1} \geq -\phi \tag{1.14}$$

Following Aiyagari (1994) the borrowing limit satisfies:

$$\phi = \left\{ \begin{array}{c} \min\left[\gamma, \frac{\omega^{\min}}{r}\right] \text{ if } r > 0 \\ \gamma \text{ if } r \leq 0 \end{array} \right\} \tag{1.15}$$

where $\omega^{\min}$ represents the worst possible income state.[13] In this case $\frac{\omega^{\min}}{r}$ represents the *natural debt limit* and γ is an *ad hoc* debt limit. The *natural debt limit* is based on the idea that banks will only lend up to the point where an indebted individual can repay their debt with certainty, even if they happen to draw the worst possible income realization at every point in the future. This *natural debt limit* is derived under the assumption, that the lender can confiscate all income received by the debtor for purposes of debt repayment. This is evidently a rather unrealistic assumption and is therefore rarely encountered in applications, where *ad hoc limits* are commonly used. However, the reasoning behind the *natural debt limit* gives us a strong theoretical reason for the existence of a finite lower bound on assets.

In line with this reasoning, the set of assets is defined as $A = [-\phi, \infty)$ with associated $\sigma-$algebra $\not{A}$.

I assume that the instantaneous utility function $u : [0, \infty) \Rightarrow \mathbb{R}$ is bounded, twice continuously differentiable as well as strictly increasing and strictly concave. Furthermore, it is assumed that the first derivative of u satisfies $\lim_{c\to\infty} u_c(c) = 0$ and $\lim_{c\to 0} u_c(c) = \infty$ i.e. the marginal utility derived from consumption approaches 0 as consumption approaches infinity, and vice versa. Finally, I assume that the degree of absolute risk aversion tends to 0 as consumption tends to infinity: $\lim_{c\to\infty} \inf \left(-\frac{u_{cc}}{u_c} = 0\right)$.

These restrictions are typical for both the partial equilibrium income fluctuation literature (c.f. Miao (2014)) as well as the literature on heterogeneous agents incomplete markets models with general equilibrium (c.f. Acikgoz (2018), Aiyagari (1994)).

It is well known (c.f. Aiyagari (1994), Chamberlain & Wilson (2000), Miao (2014)), that any solution to an income fluctuation problem with finite assets must satisfy $r < \left(\frac{1}{\beta} - 1\right)$. Otherwise, the precautionary savings motive overrides the intertemporal substitution motive and households accumulate an infinite amount of assets. Hence I assume this condition, as well as $1 + r > 0$. The latter condition ensures that households cannot grow rich, by taking on debt, while it allows for the possibility that the interest rate

[13] In our case this would be any state in which the individual has drawn $s_t = (0, ..., 0)$. In order to avoid issues with zero income, I define $\omega^{\min}$ as the income resulting from draws of s_t such that an individual that has borrowed the highest possible amount can maintain nonzero consumption by again borrowing the highest possible amount: $\omega^{\min} > r\phi$. All draws of s_t that would result in a lower labour income are assigned measure 0 by Ω.

In practice, this is not going to be an issue.

might be negative. It is easy to see from the budget constraint, that - if this condition was violated - a household could increase its consumption in period $t+1$ by taking on debt in period t: For example, let $a_t < 0$ and $1+r < 0$, then the amount of wealth carried forward to period $t+1$ is $(1+r)a_t > 0$.

The exogenous state n_t describes the occupation that the worker is matched with at time t and $s_t \in \tilde{S}$ describes the worker's skill set at t. Correspondingly n_t evolves according to the occupation transition matrix Π, and s_t according to the Markov Process $(\tilde{S}, \tilde{\mathcal{S}}, \Omega)$.

The households problem is as follows: Given initial conditions $(a_0, n_0, s_0) \in A \times N \times \tilde{S}$ and taking prices $r, \{w_m^n\}_{m \in M}^{n \in N}$ as given, the typical household chooses sequences of consumption $\{c_t\}_{t=1}^{\infty}$ and assets $\{a_{t+1}\}_{t=1}^{\infty}$ that solve the following sequential utility maximization problem:

$$V(a_0, n_0, s_0) = \sup_{\{c_t, a_{t+1}\}_{t=0}^{\infty}} E_0 \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t) \right\} \tag{1.16a}$$

subject to the budget constraint:

$$c_t + a_{t+1} = (1+r)a_t + \omega_t \tag{1.17}$$

where

$$\omega_t = \sum_{n=1}^{N} \mathbf{1}_{(n=n_t)} \sum_{m=1}^{M} (w_m^n s_{m,t}) \tag{1.18}$$

is the total labour income derived from working in occupation n at time t. Here $\mathbf{1}_{(n=n_t)}$ is an indicator function, that takes the value 1, if the worker is employed in occupation n at time t and 0 otherwise.

To rewrite the sequential problem in recursive form, define $v(a_t, n_t, s_t; \{w_m^n\}_{\forall m}^{\forall n}, r)$ as the optimum value of the objective function, starting at the joint asset, occupation and skill set state (a_t, n_t, s_t). This gives a standard dynamic programming formulation:[14]

$$v(a_t, n_t, s_t) = \max_{a_{t+1} \geq -\phi} \left\{ u(c_t) + \beta \sum_{n_{t+1}=1}^{N} \left(\int v(a_{t+1}, n_{t+1}, s_{t+1}) \Omega(s_t, ds_{t+1}) \right) \Pi(n_t, dn_{t+1}) \right\} \tag{1.19}$$

For simplicity I will from now on combine the two exogenous states n_t, s_t into a single state ψ_t with associated state space $\Psi = (N \times \tilde{S})$ and $\sigma-$algebra $\mathit{\Psi}$ and transition function $\Upsilon : \Psi \times \mathit{\Psi} \Rightarrow [0,1]$. This single exogenous state now summarizes the worker's occupation

[14]Surpressing the dependence on prices for notational simplicity.

match, as well as his skill set at time t. Let the combined state evolve according to the following Markov Process $(\Psi, \not\Psi, \Upsilon)$. Hence the problem can be stated as:

$$v(a_t, \psi_t) = \max_{a_{t+1} \geq -\phi} \{u(c_t) + \beta \int v(a_{t+1}, \psi_{t+1}) \Upsilon(\psi_t, d\psi_{t+1})\} \tag{1.20}$$

We can apply standard stochastic dynamic programming results (e.g. Theorem 9.8 in Stokey & Lucas (1989)), to conclude that $v(a_t, \psi_t)$ is strictly concave in a_t and that the policy functions: $C(a_t, \psi_t) = c_t$, $G(a_t, \psi_t) = a_{t+1}$ are continuous, single valued functions. Assumption **9.4** is met because a_t is part of the real line; **9.5** is met because $\psi_t \in \Psi$ is a countable set; **9.6** holds because $a_{t+1} \in [-\phi, (1+r)a_t + \omega_t]$; **9.7** holds by assumption; **9.10** holds by the concavity of $u(c_t)$; and **9.11** holds by the linearity of the budget constraint.

The solution to the problem can be calculated numerically by value function iteration.[15]

After solving the representative households problem, we can turn to the implied joint distribution of endogenous and exogenous variables.

Let us define a transition function $T[(a, \psi), \hat{A} \times \hat{B}] : (A \times \Psi) \times (\not A \times \not\Psi) \Rightarrow [0, 1]$ for any $(a, \psi) \in \Lambda$, $\hat{A} \in \not A$ and $\hat{B} \in \not\Psi$ to be the transition function induced by the policy function $G(a_t, \psi_t)$ and the stochastic process for Ψ, with the following form:[16]

$$T[(a, \psi), \hat{A} \times \hat{B}] = \left\{ \begin{array}{c} \Upsilon(\psi, \hat{B}), \ if \ G(a_t, \psi_t) \in \hat{A} \\ 0, \ \text{otherwise} \end{array} \right\} \tag{1.21}$$

Let $\chi_t(\hat{A} \times \hat{B})$ be the distribution of households over the joint state space $\Lambda = A \times \Psi$ at time t. Through its definition, T provides a law of motion for the evolution of the joint state χ_t :

$$\chi_{t+1} = T\chi_t \tag{1.22}$$

Under the assumptions made above,[17] Acikgoz (2018) shows that the Markov process on Λ with transition function T guarantees the existence of a unique, stationary distrib-

[15] However in the practical application I solve the model with the method of endogenous gridpoints (c.f. Carroll (2006)).

[16] See Stokey & Lucas Chapter 9.6 for details.

[17] Note that this specification implies, that it is possible for the agent to rank each $\psi_t \in \Psi$ as is required for the proof to go through: For any $\tilde{\psi} \in \Psi$ we can calculate a fixed quantity of labour income via the labour income equation (1.18). Applying (1.18) to all elements of Ψ generates a mapping from the elements of Ψ and the set of wages $\{w_m^n\}_{\forall m}^{\forall n}$ to a set $\omega^{\tilde{\psi}}$ that contains all the corresponding labour incomes. Remember, that we have defined the smallest element of this set $\omega^{\min} > 0$ when discussing the borrowing constraint. Furthermore, given the assumptions made about Ψ and λ, we also know that $\omega^{\tilde{\psi}}$ has a finite maximum $\omega^{\max}$, which is all that is needed.

ution $\bar{\chi}(A \times \Psi)$ that satisfies:

$$\bar{\chi} = T\bar{\chi} \tag{1.23}$$

It is further guaranteed, that the expected value of the assets held, using the invariant distribution, is continuous in r.

4 General Equilibrium

The law of motion T can be interpreted as describing the evolution of the joint endogenous and exogenous states of a typical household over time. And as the distribution of states that are visited by the typical household over an infinitely long period of time. In line with the literature (c.f. Aiyagari (1994), Ljungqvist & Sargent (2012)) I interpret χ_t as a cross-sectional distribution of households over the joint state space Λ. In particular, I interpret the measure $\chi_t(\hat{A} \times \hat{B})$ as the fraction of households whose asset positions and exogenous states lie within the set $\hat{A} \times \hat{B}$ at time t.

Aggregation over this cross-sectional distribution of households follows from a Strong Law of Large Numbers (Acemoglu & Jensen (2015)). In particular, following the concept of an *aggregator* described in Acemoglu & Jensen (2015) individual-level uncertainty can be cancelled out by integrating over the distribution of households. This implies that all aggregate quantities are fixed and therefore non-random from the perspective of the individual household.

In particular, the total capital supply is given by:

$$K = \int_{i \in I} a_{i,t} di \tag{1.24}$$

and the total labour supply by:

$$L = \sum_{n=1}^{N} \bar{\mu}_n \sum_{m=1}^{M} \lambda_m^n \bar{\eta}_m = \sum_{n=1}^{N} \bar{\mu}_n \sum_{m=1}^{M} \lambda_m^n \int_{i \in I} s_{i,m} \Omega(d\bar{s}) \tag{1.25}$$

Using these stationary quantities, I define a *competitive, recursive, stationary equilibrium* as follows:[18]

A value function $v(a_t, \psi_t) : \Lambda \Rightarrow \mathbb{R}$, and policy functions $C(a_t, \psi_t) : \Lambda \Rightarrow \mathbb{R}_+$ and $G(a_t, \psi_t) : \Lambda \Rightarrow A$; similarly an invariant distribution $\bar{\chi}(\hat{A} \times \hat{B})$ and a law of motion T; an aggregate stock of capital K, a labour aggregate L; a set of wage rates $\{w_m^n(K)\}_{m \in M}^{n \in N}$ and an interest rate $r(K)$ such that:

[18] See for example Acikgoz (2018), Ljungqvist & Sargent (2012, ch. 18), Miao (2014, ch. 17).

1. The representative firm maximizes profits so that the set of $\{w_m^n\}_{\forall m}^{\forall n}$ and r solve the firm's problem (1.4). Specifically, the set of wages $\{w_m^n\}_{\forall m}^{\forall n}$ solves the equations described by (1.7a) and the interest rate r solves (1.8).

2. Given the aggregate quantities, K and L, and prices r, $\{w_m^n\}_{\forall m}^{\forall n}$, C and G solve the household's problem (1.20) and v solves the Bellman equation (1.20).

3. T is the law of motion induced by G and Υ and $\bar{\chi}$ is the stationary distribution defined by: $\bar{\chi} = T\bar{\chi}$.

4. The capital market clears: $K = \int\limits_{i \in I} a_{i,t} di$.

Following standard arguments (e.g. Aiyagari (1994)) it can be shown that under the conditions set out above, a general equilibrium exists for this economy. Uniqueness of the stationary equilibrium is established in Acikgoz (2018).

5 Simple Example

To aid intuition I will provide a simple example of the productivity process in this section. Suppose there is only one skill h and two occupations n^{high}, n^{low} differentiated by how intensely they use the skill h. Specifically, n^{high} is associated with λ^{high} and n^{low} is associated with λ^{low}. The exogenous occupation transition matrix is given by:

	n^{high}	n^{low}
n^{high}	π_{hh}	π_{hl}
n^{low}	π_{lh}	π_{ll}

with an associated stationary distribution $\bar{\mu} = (\bar{\mu}^{high}, \bar{\mu}^{low})$.

The skill h can take on two values: h^{high}, h^{low} with transition probabilities:

	h^{high}	h^{low}
h^{high}	ϖ_{hh}	ϖ_{hl}
h^{low}	ϖ_{lh}	ϖ_{ll}

with an associated stationary distribution $\bar{s} = (\bar{s}^{high}, \bar{s}^{low})$.

The joint state space Ψ is given by the set
$\{(\lambda^{high}, h^{high}), (\lambda^{high}, h^{low}), (\lambda^{low}, h^{high}), (\lambda^{low}, h^{low})\}$ with transition matrix Υ:

	$(\lambda^{high}, h^{high})$	$(\lambda^{high}, h^{low})$	$(\lambda^{low}, h^{high})$	(λ^{low}, h^{low})
$(\lambda^{high}, h^{high})$	$\pi_{hh}\varpi_{hh}$	$\pi_{hh}\varpi_{hl}$	$\pi_{hl}\varpi_{hh}$	$\pi_{hl}\varpi_{hl}$
$(\lambda^{high}, h^{low})$	$\pi_{hh}\varpi_{lh}$	$\pi_{hh}\varpi_{ll}$	$\pi_{hl}\varpi_{lh}$	$\pi_{hl}\varpi_{ll}$
$(\lambda^{low}, h^{high})$	$\pi_{lh}\varpi_{hh}$	$\pi_{lh}\varpi_{hl}$	$\pi_{ll}\varpi_{hh}$	$\pi_{ll}\varpi_{hl}$
(λ^{low}, h^{low})	$\pi_{lh}\varpi_{lh}$	$\pi_{lh}\varpi_{ll}$	$\pi_{ll}\varpi_{lh}$	$\pi_{ll}\varpi_{ll}$

and associated stationary distribution $\varrho = (\bar{\mu}^{high}\bar{s}^{high}, \bar{\mu}^{high}\bar{s}^{low}, \bar{\mu}^{low}\bar{s}^{high}, \bar{\mu}^{low}\bar{s}^{low})$. The labour supply in the first occupation l^{high} is therefore given by:

$$l^{high} = \left(\left(\bar{\mu}^{high}\bar{s}^{high} * \lambda^{high} * h^{high}\right) + \left(\bar{\mu}^{high}\bar{s}^{low} * \lambda^{high} * h^{low}\right)\right)$$

with l^{low} similarly defined:

$$l^{low} = \left(\left(\bar{\mu}^{low}\bar{s}^{high} * \lambda^{low} * h^{high}\right) + \left(\bar{\mu}^{low}\bar{s}^{low} * \lambda^{low} * h^{low}\right)\right)$$

Total labour supply is given by:

$$L = l^{high} + l^{low} \tag{1.26}$$

Given L one can easily derive the set of occupation-specific wage rates w^{high} and w^{low} for any value of the aggregate capital stock K:

$$w^{high} = \frac{(1-\alpha)\lambda^{high}K^{\alpha}}{L^{\alpha}} \tag{1.27}$$

$$w^{low} = \frac{(1-\alpha)\lambda^{low}K^{\alpha}}{L^{\alpha}} \tag{1.28}$$

Taking these as given the worker knows his labour income in each of the 4 states $\psi \in \Psi$. Taking these as given one can solve the typical household's problem and the rest of the model in a familiar fashion.

After illustrating how the aggregate Labour Supply is derived in this setting let us turn to the properties of the implied income process. By comparing the *structural* income process defined by the model with a reduced-form version, I will highlight some of the benefits derived from thinking about the structure of income determination in the given task-skill framework. Most HIM models can be calibrated using information about the persistence and unconditional variance of income. For simplicity, I will focus on the unconditional variance, but similar arguments can be made using the persistence.

Average labour income in this economy is given by:

$$E(\omega) = \begin{bmatrix} (\bar{\mu}^{high}\bar{s}^{high}w^{high}h^{high}) + (\bar{\mu}^{high}\bar{s}^{low}w^{high}h^{low}) \\ +(\bar{\mu}^{low}\bar{s}^{high}w^{low}h^{high}) + (\bar{\mu}^{low}\bar{s}^{low}w^{low}h^{low}) \end{bmatrix} \tag{1.29}$$

The variance of labour income is given by:

$$Var(\omega) = \begin{bmatrix} \bar{\mu}^{high}\bar{s}^{high} * (w^{high}h^{high})^2 + \bar{\mu}^{high}\bar{s}^{low} * (w^{high}h^{low})^2 \\ +\bar{\mu}^{low}\bar{s}^{high} * (w^{low}h^{high})^2 + \bar{\mu}^{low}\bar{s}^{low} * (w^{low}h^{low})^2 \end{bmatrix} - E(\omega)^2 \tag{1.30}$$

Now suppose that the econometrician is unable to observe changes to an individual's skill-set or occupation match, but instead only observes total labour income $y = w * h$ in the 4 different income states. Denote these states (in order of appearance above) by $y_1, .., y_4$ and $p_{11}, ..., p_{44}$ as the associated transition probabilities between the states.[19]

Together, $y = \{y_1, .., y_4\}$ and the transition matrix P defined by $p_{11}, ..., p_{44}$ describe a reduced form income process that will be familiar from many standard HIM applications that utilize discrete state Markov chain approximations to autoregressive income processes (c.f. Aiyagari (1994), Rouwenhorst (1995), Tauchen (1986)).

In particular, letting $\bar{p}_n$ denote the unconditional probability of being in income state $n = \{1, .., 4\}$.

Average labour income given by the reduced form approach is:

$$E(y) = (\bar{p}_1 y_1 + \bar{p}_2 y_2 + \bar{p}_3 y_3 + \bar{p}_4 y_4) \tag{1.31}$$

And the variance is given by:

$$Var(y) = \left(\bar{p}_1 y_1^2 + \bar{p}_2 y_2^2 + \bar{p}_3 y_3^2 + \bar{p}_4 y_4^2\right) - E(y)^2 \tag{1.32}$$

Comparing equation (1.29) with (1.31) and (1.30) with (1.32) we can spot the problem, that this approach is trying to address. If we are able to observe the reduced-form labour income states, we should expect both approaches to deliver the same results. However, if we were trying to anticipate changes to the two moments of the income distribution, we would struggle to find a solid basis for changes to the parameters of reduced-form income process. Without grounding the determinants of income in any structural relationship, forecasting the evolution of any moment of the income distribution will involve a degree of arbitrariness.

[19] Every p is a combination of π and ϖ.

Focussing on the determination of labour income through $y = w(\lambda) * h$[20] we can see that changes to observed labour income y can be driven by both, technological factors - as embodied in $w(\lambda)$ - and changes in the distribution of human capital. The former effect might be dominant throughout a period of rapid technological change, whilst the latter could be the driving force after a major educational reform.

The same can be said of the transitions between the different income states $p = \pi * \varpi$ where π might represent structural behaviour in the labour market and ϖ the speed at which individuals change their skill set. In both cases, the two principal actors of Tinbergen's *race* interact to produce one indivisible outcome.

The rest of the chapter essentially focuses on addressing this issue. Exploiting the fact that the structural and reduced-form income processes should produce the same results, I am able to calibrate the structural model with observable data. Then I can use changes to the structural environment to generate changes in the income process that are based on popular theories on technological change in order to observe the effects on income and wealth inequality.

The next section describes the calibration of the structural model in a stationary environment, setting the stage for the policy experiments in the following sections.

6 Calibration

6.1 Introduction

I calibrate a stationary version of the model using data from 8 waves of the Understanding Society (UnSoc.) survey, covering the years 2009 - 2018. Unless differently stated the analysis focuses on white British males, aged between 25 and 55, who are employed for a number of consecutive periods.[21]

The sample is selected to represent the closest approximation to a truly competitive labour market, in which observed pay is most closely related to actual productivity. This is in line with much of the literature on labour market outcomes. As a measure of income, I use the usual net pay per month in the current job, deflated by the CPI. In order to avoid issues with top (bottom) coded incomes I trim the top and bottom 0.5% in terms of labour income.[22] As with any survey, Understanding Society contains a number of accidental errors that occur along the way from the interview to the finished dataset. For

[20]Here I made the dependence of w on λ explicit.

[21]So if I observe some individual as employed for a number of periods who becomes unemployed after a certain period I drop all further observations on this individual.

[22]Choosing 0.25% as cutoffs does not significantly alter the results. As does not excluding any observations.

example, a number of individuals are reported to earn £1 per month which can be safely attributed to data error. By removing the worst outliers I ensure that the estimates are not disturbed by too much random noise.

6.2 Skill Measures & Processes

For this application, I limit myself to two types of general skills: Cognitive and Physical.[23] These can be seen as embodiments of the two pillars of human capital theory, with the former representing intellectual capacity, and the latter bodily health and physical ability. Apart from this, the distinction between analytic and manual tasks has long been at the heart of the literature on task-biased technological change (c.f. Autor et al. (2003)). In line with the Task-Skill framework, I suggest that cognitive skills will be applied to cognitive tasks - i.e. there is a skill *Cognitive* amongst the $M = 2$ skill types, and a set of corresponding cognitive skill weights: $\left\{\lambda^n_{Cog}\right\}^N_{n=1} \in \lambda$. And similarly for Physical skills.

Using test scores for a variety of numerical and verbal intelligence tests in wave 3 of the UnSoc, I generate a composite measure of cognitive ability by performing a Principal Component Analysis and selecting the first principal component. The method here follows Whitley et al. (2016), who use the same survey to study the effect of ageing on cognitive decline.[24]

The UnSoc surveys also include the SF-12 Physical Component Summary (PCS) Index, which is a composite index evaluating the physical health of an individual (c.f. Ware et al. (2001)).[25] This is a self-completed questionnaire that encompasses 12 questions regarding different aspects of physical health and fitness. The PCS score is then derived from the answers, providing a summary statistic rating the respondent's health. For lack of alternatives, I take this as a proxy for an individual's physical ability (c.f. Lise & Postal-Viney (2017)). I standardize both scores on the interval [0,1] in order to make them comparable.

The cross-sectional correlation between both measures is around 0.10, so I believe that I am justified in treating both processes as independent.

For simplicity, the evolution of both skills are defined as stationary AR(1) processes in logarithms:

[23] At this point it appears useful to remind ourselves, that there are many possible ways of dividing the complex manifold of Tasks and Skills. One could have split the workers into "skilled" and "unskilled" workers, potentially employing some sort of educational qualification as dividing characteristics. Alternatively one could have employed the complex-routine/manual-cognitive distinction first used by (Autor et al. (2003)). Ultimately the final choice comes down to preference and data availability.

[24] For more information on this procedure see the Appendix.

[25] Guvenen et al. (2017) use this index to proxy physical ability.

$$\log(s_{m,t+1}) = \rho_m \log(s_{m,t}) + \varepsilon^m_{t+1} \tag{1.33a}$$

$$|\rho_m| < 1 \tag{1.33b}$$

$$\varepsilon^m_t \sim N(0, \sigma^2_m) \tag{1.33c}$$

Their calibration will provide an important part of the income dynamics faced by the workers. I use the availability of multiple measures of the SF-12 PCS in the data (the test is completed every year) to estimate the following regression equation:

$$\log(s^{Phy}_{t+1}) = \rho_{Phy} \log(s^{Phy}_t) + \varepsilon^{Phy}_{t+1} \tag{1.34a}$$

$$\varepsilon^{Phy}_t \sim N(0, \sigma^2_{Phy}) \tag{1.34b}$$

Unfortunately, a similar approach is not possible for the measure of cognitive ability, since I only have a single observation per individual.[26] I therefore, chose the relevant parameters to match key moments in the data.[27]

The table below summarizes the relevant parameters:

	Cognitive	Physical
$\overline{\eta}$	0.681	0.685
ρ	0.867^+	0.909
σ_ε	0.245^+	0.154

Note: A + indicates the parameter is calibrated.

Table 1: Skill Processes

The pattern that emerges from these values seems reasonably intuitive. Both processes have a high level of persistence with a standard deviation of innovations between 15 and 25%. Unfortunately, there is - to the best of my knowledge - no study that approaches the evolution of skills in a similar manner, and therefore it is impossible to verify whether these calibrations are reasonable. The closest comparison study might be Huggett et al. (2011) who estimate a standard deviation of shocks to human capital of around 11%. Since shocks to human capital in this model are composed of shocks to either skill, as well as occupation transitions, the values do not seem too far off the mark.

[26] According to USoc, another cognitive skills test is planned in wave 9, which might provide a possibility for checking the validity of my calibration.

[27] The target moments are provided by targeting the reduced form log-normal income process in the data, namely the persistence of the AR(1) process, the standard deviation of innovations and the standard deviation of log(Income). For the calibration procedure, see Appendix.

For the computational implementation I approximate each (1.33a) by a 5-state discrete Markov Chain using the Rouwenhorst (1995) method.[28]

6.3 Skill Weights

I now explain how to recover the skill weights λ from the data. Many recent papers attempt to estimate the return to specific work activities - usually referred to as "task returns", which play a similar role as the skill weights referred to here - by using survey information on the importance of certain work tasks (c.f. Autor & Handel (2013), Baumgarten et al. (2013), Firpo et al. (2011)). These approaches usually have to deal with selection bias, where individuals with high unobserved ability in certain tasks, sort into occupations where this task is used very intensely according to a Roy framework (c.f. Autor & Handel (2013), Firpo et al. (2011), Roys & Taber (2016)). This means that if skill levels are unobserved, estimated task returns will be biased upwards under positive sorting. The main benefit of the route taken here is that skills are observed and can thus be controlled for.[29]

Recall that an individual i, in occupation n, earns labour income $\omega_{i,n}$[30] equal to:

$$\omega_{i,n} = \sum_{m=1}^{M} (w_m^n s_{i,m}) \tag{1.35}$$

substituting the wage equation (1.7a) into this equation gives:

$$\omega_{i,n} = \sum_{m=1}^{M} \left(\frac{(1-\alpha)\lambda_m^n K^\alpha}{L^\alpha} \right) s_{i,m} \tag{1.36}$$

noting that some of these terms are invariant across all m we can rewrite this equation as:

$$\omega_{i,n} = \underbrace{(1-\alpha)\left(\frac{K}{L}\right)^\alpha}_{\text{common}} \underbrace{\sum_{m=1}^{M} \lambda_m^n s_{i,m}}_{\text{individual}} \tag{1.37}$$

It is easy to see that the above equation can be split into a common part $(1-\alpha)\left(\frac{K}{L}\right)^\alpha$ that is constant across all individuals in the economy, and a worker-specific part $\sum_{m=1}^{M} \lambda_m^n s_{i,m}$ that depends on the worker's individual skill set and occupation match.

[28] The discretization error resulting from this procedure is of the order of 5%.

[29] Regressing wages on skills is the procedure suggested by Autor (2013).

[30] For expositional clarity, I make the dependence of income on the occupation match n explicit.

Our estimation will focus on the individual-specific part. For the econometric specification, I will treat the common part as a constant $\Xi = (1-\alpha)\left(\frac{K}{L}\right)^{\alpha}$ that multiplies each $\lambda_m^n \in \lambda$. Hence I will redefine the skill weights to include the level information included in the common constant:

$$\tilde{\lambda}_m^n = \Xi * \lambda_m^n \tag{1.38}$$

Using this transformation we derive the following expression:

$$\omega_{i,n} = \sum_{m=1}^{M} \tilde{\lambda}_m^n s_{i,m} \tag{1.39}$$

Since $\omega_{i,n}$ and $s_{i,m}$ are observed in the data, I can estimate (1.39) for a given n, and obtain estimates of $\tilde{\lambda}^n$ as the coefficients of a simple OLS regression. Furthermore, I can pool all n and use dummy variables to estimate the whole set of $\tilde{\lambda}$ jointly:

$$\omega_{i,n} = \sum_{n=1}^{N} \left\{ D_n \sum_{m=1}^{M} \tilde{\lambda}_m^n s_{i,m} \right\} \tag{1.40}$$

Here D_n is an occupation dummy that is equal to 1 if the individual is employed in occupation n and zero otherwise. I obtain estimates of the rescaled skill weights $\tilde{\lambda}_m^n$ as the coefficients of the measured skills $s_{i,m}$ interacted with the occupation dummies D_n. This econometric strategy draws on Deming (2017) and Gensowski (2017) who investigate the return to different skills using similar reduced-form specifications.

The results of this regression for all 9 major SOC2000 occupations are reported below:

In order to reduce the number of variables, I separate the 9 major occupational groups into three larger occupation clusters. For this, I run a hierarchical clustering algorithm on the skill weights obtained from this initial regression. The clustering algorithm groups the occupations according to their similarity with respect to the estimated skill weights. The resulting clusters[31] loosely represent some of the key worker types in the skill and task-biased technological change literature:

(1) Managers, Professionals & Associate Professionals, representing traditional high-skilled occupations that are intensive in complex analytic tasks.

(2) Administrative & Sales occupations, represent middle to low-skilled occupations with a focus on routine cognitive tasks.

Finally (3), Skilled Trades, Personal Services, Plant & Machine Operatives & Miscellaneous occupations, cover workers that work in manual tasks.[32]

[31] I use single linkage, average linkage and Ward's linkage, all of which result in the reported clustering.

[32] In order to save space I will sometimes refer to these groups by abbreviations; e.g. "Admin" for

	Labour Income
Managers X Cognitive Skills	2136.867***
	(237.272)
Professionals X Cognitive Skills	2012.366***
	(305.273)
Associate Professional X Cognitive Skills	2169.107***
	(242.370)
Administrative X Cognitive Skills	1635.570***
	(259.414)
Skilled Trades X Cognitive Skills	1358.362***
	(207.406)
Personal Services X Cognitive Skills	1096.985***
	(308.619)
Sales X Cognitive Skills	1407.666***
	(334.081)
Plant & Machine Operatives X Cognitive Skills	1338.344***
	(159.251)
Misc X Cognitive Skills	998.316***
	(148.150)
Managers X Physical Skills	1327.115***
	(239.439)
Professionals X Physical Skills	1416.625***
	(314.107)
Associate Professional X Physical Skills	799.192***
	(241.528)
Administrative X Physical Skills	814.155**
	(261.155)
Skilled Trades X Physical Skillsl	1336.246***
	(194.737)
Personal Services X Physical Skills	1072.406***
	(296.482)
Sales X Physical Skills	530.833
	(288.284)
Plant & Machine Operatives X Physical Skills	1254.784***
	(146.051)
Misc X Physical Skills	1109.200***
	(134.661)
R-squared	0.878
N	4393

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Skill Weights Regression - Major SOC2000 Occupations

The clear interpretability of the different clusters already suggests, that the occupation-specific skill weights $\tilde{\lambda}^n$ that I have obtained from the regression pick up information about the structural differences between the different occupations. To analyze this further, I repeat the regression analysis with the reduced set of occupations.

	Labour Income
Managers & Professionals X Cognitive Skills	2134.534***
	(151.876)
Admin & Sales X Cognitive Skills	1692.235***
	(214.909)
Skilled Workers, Services, Operatives & Misc X Cognitive Skills	1273.780***
	(97.094)
Managers & Professionals X Physical Skills	1172.172***
	(153.777)
Admin & Sales X Physical Skills	593.662**
	(207.421)
Skilled Trades, Services, Operatives & Misc X Physical Skills	1217.692***
	(89.798)
R-squared	0.873
N	4393

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Skill Weights Regression - Occupation Clusters

The regression table shows clear, discernible differences in the returns to cognitive and physical skills across the three occupation clusters. In order to aid further interpretation, I rescale the estimated skill weights, so that the implied total labour supply is standardized to 1. The results are shown below:

	Managers & Professionals	Admin & Sales	Trades, Services, Operators & Misc.
$\lambda^{Cognitive}$	1.04	0.82	0.62
$\lambda^{Physical}$	0.57	0.29	0.59
$\frac{\lambda^{Cognitive}}{\lambda^{Physical}}$	1.82	2.85	1.05

Table 4: Standardized Skill Weights

Taken at face value the results are quite revealing by themselves: Cognitive skills seem to be associated with higher productivities through the board, with a reasonable ranking of the cognitive skill weights across the different occupations. Managers, Professionals and Associate Professionals tend to have higher returns to their cognitive skills than the

"Administrative & Sales".

other occupations, with Skilled Workers, Services, Machine Operators & Misc. exhibiting a sizeable 40% gap to the Managers & Professionals group.

Curiously Managers, Professionals and Associate Professionals also have relatively high returns to Physical skills, which however might reflect some higher general level of productivity that is unobserved. My preferred interpretation of this observation is that Physical skills are less proxies for physical strength, and more indicative of stamina and resilience. Given that many managerial and professional occupations can be very demanding and stressful, it should not be surprising that the returns to physical health are high.

Setting this issue aside for the moment, I also confirm that Skilled workers have high returns to Physical skills - in particular when compared to their cognitive returns. Administrative & Services meanwhile have the lowest returns to Physical skills both in absolute and in relative terms.

The ratio of the cognitive and the physical skill weight provides us with an indication of the degree of skill risk faced by each occupation: Since skill shocks are independent, occupations with a high ratio are much more exposed to fluctuations in labour income due to high or low realizations of the cognitive skill component. Administrative and Sales workers for example seem to be very specialized in cognitive tasks, with a skill weight over three times higher than the physical skill weight. This implies that they are much more exposed to changes in their cognitive skills and, correspondingly face higher within-group income inequality. This stands in stark contrast to the more balanced Skilled worker types, whose occupations place about the same weight onto both types of work activities. Correspondingly we should expect these to be less exposed to income risk due to skill fluctuations and also exhibit lower within-group income inequality.

Overall the results seem to confirm a picture where the returns to different skills are differentiated across different occupations, with a somewhat clear pattern of task specialization.

6.4 Occupation Transitions

I estimate the yearly[33] job-to-job transition probabilities from the data.

[33] Technically these are wave-wave transitions, but in practice, these almost perfectly line up as USoc tries to interview with a yearly rhythm.

	Managers	Admin	Trades
Managers	97.92	0.79	1.29
Admin	8.21	88.12	3.67
Trades	2.92	0.86	96.22
Total	58.22	9.39	32.39

Table 5: Cluster Transition Probabilities (Percentages)

As can be seen from the estimated occupation transition matrix, transitions between occupational clusters are relatively rare. Even for the Administrative & Sales cluster, the probability to stay in the same cluster year by year is 88%. The low turnover can probably be attributed to two factors:

1. It is well known, that individuals are more likely to change their occupation after an unemployment spell (c.f. Carrillo-Tudela et al. (2016)), however, I have excluded unemployed individuals from the sample.

2. Gathmann & Schoenberg (2010) show, that occupation transitions are more likely to occur between occupations that have similar task content. Since the occupational clusters have been constructed based on the similarity of task content, it is very likely that the level of aggregation masks occupation transitions occurring between the component occupational groups of each cluster.

6.5 Other Parameters

I set the functional form of the utility function to *CRRA* and calibrate the coefficient of relative risk aversion $\varkappa = 2$, and the time preference $\beta = 0.975$, which yields a risk-free interest rate of 2.56%. In line with most calibrations, I set $\alpha = 0.36$ and the rate of depreciation to $\delta = 0.1$. I set the ad hoc borrowing limit $\gamma = 0$. For the final calibration, I standardize the occupation-specific productivities so that the total Labour Supply is equal to 1.

7 The Stationary Model

7.1 Income - Overall Sample

Since the aim of this model is to put some structure on the earnings process, it will be important to evaluate the model-generated income process against the dynamics found in the data. As a benchmark, I use a reduced form earnings process of the form:

$$\log(y_{i,t+1}) = \kappa \log(y_{i,t}) + \varepsilon_{i,t+1} \tag{1.41a}$$

$$|\kappa| < 1 \tag{1.41b}$$

$$\varepsilon_{i,t} \sim N(0, \sigma_\varepsilon^2) \tag{1.41c}$$

I estimate the parameters for κ and σ_ε from the UnSoc data and use them to calibrate the productivity process in the model. After solving the model I simulate 5000 individual labour market histories, each for 40 periods in order to obtain the relevant parameter estimates. The table below shows the estimated parameters of the income process for the data and the models:

	κ	σ_ε	σ_y	Gini	$\frac{Mean}{Median}$
Data	0.882	0.155	0.399	0.23	1.09
Model	0.879	0.173	0.388	0.22	1.09

Table 6: Income Process

Comparing the parameter estimates from the model with the ones from the data immediately highlights what a good job the model does at reproducing the patterns in the data. The persistence and the standard deviation of income idiosyncratic income shocks are fairly well matched. Across the cross-sectional dimension, the model matches the data closely.

7.2 Income - Within and between Occupations

A main advantage of the model is that we can attribute meaningful interpretations to the different productivity states: namely, they correspond to employment in different occupations. Within- and between occupation income inequality has recently become a topic of interest for macroeconomists (c.f. Kambourov & Manovski (2009)). The following table compares the within- and between-occupation income inequality produced by the model with that observed in the data.[34]

It can be seen that the model is able to capture the relative ranking of the different occupations by their mean income fairly well and does a decent job at capturing some of the between-group inequality observed in the data. It only slightly overshoots with respect to the mean income of the Skilled, Services, Operatives & Misc. cluster.

[34] Here y is the logarithm of income. A bar indicates the average.

	$\bar{y}_{Data}$	$\bar{y}_{Model}$	$\bar{y}_{Data} - \bar{y}_{Model}$	$\frac{\sigma^y_{Data}}{\sigma^y_{Model}}$	$Gini_{Data}$	$Gini_{Model}$
Managers	0.16	0.13	0.03	1.04	0.20	0.20
Admin	-0.23	-0.24	0.02	0.86	0.20	0.23
Trades	-0.19	-0.14	-0.05	1.02	0.19	0.18

Table 7: Between and Within Occupation Income Inequality

Further, the model does a decent job with respect to within-group inequality, both measured by the within-group standard deviation of y as well as the Gini coefficient. In general, however, the between-group inequality tends to be between a little higher in the data than what the model can reproduce. This should not be too surprising, as the model does not allow for selection into preferred occupations, or account for any source of ex-ante skill heterogeneity. Nonetheless, the model produces a surprisingly accurate picture of between- and within occupation income inequality, suggesting that the model has some merit.[35]

7.3 Assets

Having established that the model is able to reproduce important features of the data with respect to the income process, let us turn to the quantitative predictions of the model with respect to assets. Unfortunately, the UnSoc. survey does not contain much information on asset holdings so that a comparison with the empirical distribution is not possible. Instead, I compare the results to some measures of wealth inequality derived from the Wealth and Assets (WAS) survey,[36] as well as a standard Aiyagari (1994) model that has been calibrated using a 12-state discrete Markov chain representation of the income process (1.41a).[37] This has long been an established approach in the literature and should provide us with a solid comparison for the model's predictions.

The table below summarizes some important features of the asset distribution implied by the model, the WAS, as well as the baseline Aiyagari:

Quantitatively the predictions borne out by the baseline Aiyagari model and the one with skills and occupations appear to be quite similar. There appears to be a slightly higher level of savings - and conversely, lower level of wealth inequality - in the baseline model, but the differences are only minor. It is reassuring to see the similarity between both models, as it suggests that the model does not generate any "unusual" asset dynam-

[35] In particular when one considers that these moments were not directly targeted by the calibration.
[36] Waves 3 & 4.
[37] For this approximation I use Rouwenhorst (1995).

	Model		Standard Aiyagari		WAS*	
	Mean	Gini	Mean	Gini	Mean	Gini
Overall	1	0.51	1.03	0.50	1**	0.62
Managers	1.14	0.50	-	-	1.41	0.57
Admin	0.81	0.58	-	-	0.75	0.66
Trades	0.73	0.58	-	-	0.53	0.63

Note:* total personal wealth ** standardized to allow relative group comparisions

Table 8: Wealth Inequality

ics.

Overall it is notable that the implied wealth inequality is far below empirical estimates. This is particularly evident when looking at the values of the Gini coefficient: the values here range around 0.5, whilst the empirical estimate is 0.62. It is important to emphasize that this is potentially due to the fact that the income process does not take into account all sources of income variation but rather focuses on income inequality due to differences in human capital. However, the model does a decent job at capturing the between- and within-group wealth inequalities, even though it doesn't quite reproduce quite so extreme disparities, as the WAS suggests. Managers hold, on average about 50% more wealth than the other occupational groups, which have a similar level of mean wealth. Conversely, within group wealth inequality is about 8 Gini points higher within the latter two groups.

The graph below shows the stationary distribution of assets. As expected, the patterns described by both models are almost indistinguishable. The takeaway from this exercise is that the model's implications for wealth inequality are very close to those obtained from the standard model. Having established the properties of the stationary model, I will use it to perform several pieces of policy-relevant analysis in the following sections.

Figure 1: Stationary Distribution of Assets

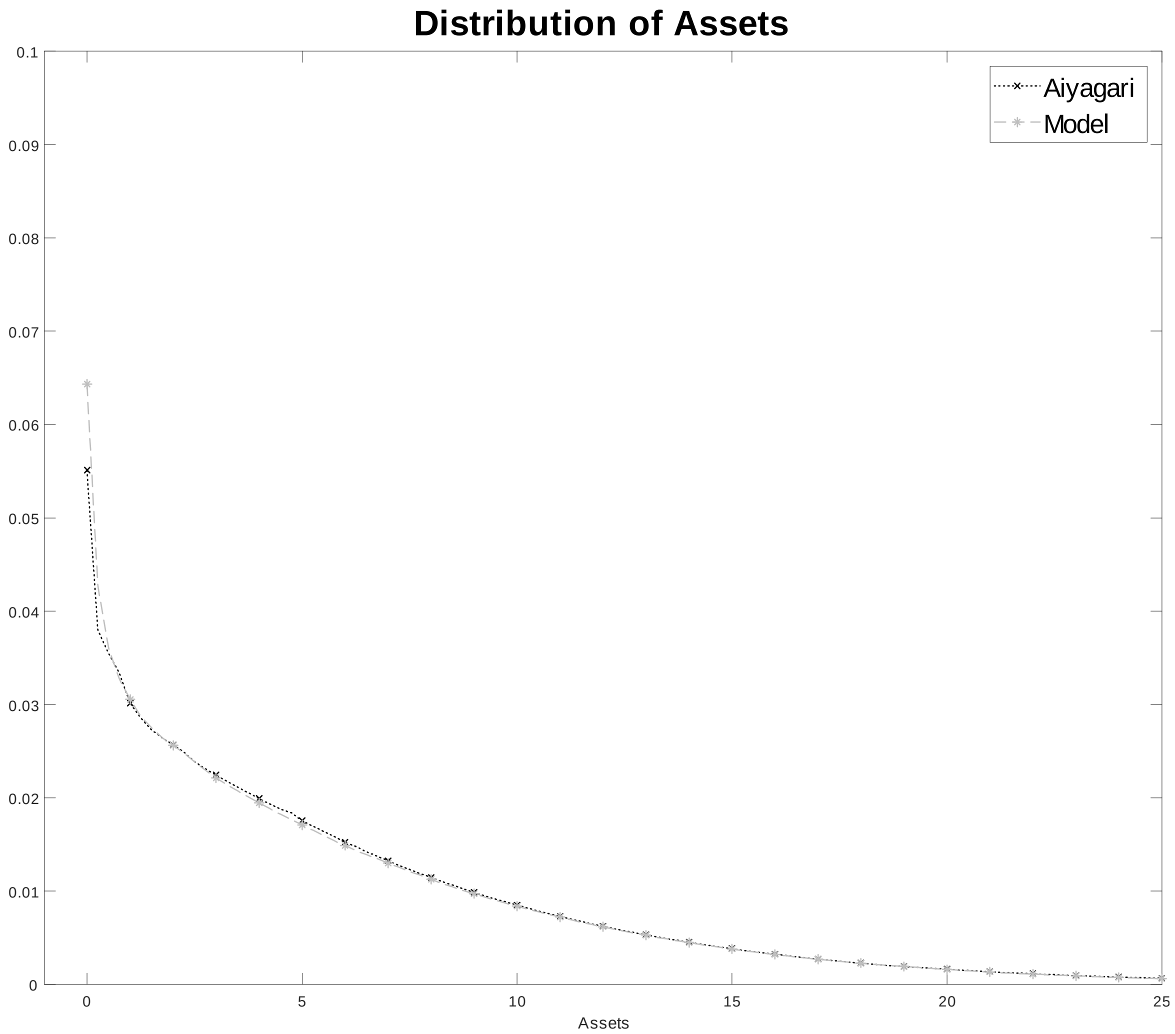

Note: Comparison of model with standard Aiyagari

8 Income Risk Decomposition

The key difference between the economic model presented in the last sections and the standard Aiyagari (1994) model, is that the former decomposes the variation of an individual worker's income into two parts: 1. Variation in the amount of skills that the worker can supply, and 2. Changes in the demand for the skills of the individual worker due to occupational transitions. Both aspects are relevant to understanding income fluctuations but are representing fundamentally different parts of economic life.

In the following section, I will utilize this decomposition, to analyze the effect of each aspect on wealth & consumption inequality.

The skill supply in this model is governed by the autoregressive skill process (1.33a), which determines the stock of cognitive and physical skills available to the individual worker. My preferred interpretation is, that shocks to the individual's skillset represent life events that are at least somewhat exogenous. This is quite evident with respect to physical skills, where accidents and illness might play the role of exogenous shocks, but similarly, a sudden training opportunity might be reflected in a positive cognitive skill shock.

Adopting this interpretation, the standard deviation of skill shocks becomes a measure of inequality of opportunity. A high spread of shocks to physical skills for example suggests that some individuals have better access to healthcare facilities than others, whilst a low spread of innovations to cognitive skills would e.g. suggest an approximately equitable distribution of educational opportunities.

Similarly, the set of skill weights summarizes the productivity (and thus demand for) differentiated skills. As technology evolves these values are liable to change, posing a source of considerable risk for workers, who might find that their skills are no longer in demand.[38]

To assess the differentiated impact of these different sources of risk might be of interest to a variety of policymakers, and being able to provide an analysis at this disaggregated level is one of the key advantages of the model presented here.

In particular, a government might try to enact legislation in order to affect the risk posed by these different channels. These could be reforms to the healthcare or educational system, acting as insurance against physical and cognitive shocks respectively. Similarly, a government might try to encourage some type of technological progress, by imposing suitable regulations on different industries,[39] thereby affecting the demand for different skills.

The structure of the model provides an opportunity, to evaluate the impact of these different policies on the mean and distribution of income, consumption and wealth.

The following tables present a set of comparative statics exercises, investigating the effect of changes to the risk posed by cognitive and physical skill shocks, as well as changing technological demands.

Holding all other parameters constant, I vary the size of σ_{Cog}, σ_{Phy}, λ^{Cog} & λ^{Phy} and solve the model for each specification, obtaining the stationary distribution of labour

[38] More on this topic in the next section.

[39] See for example the suggestions in Atkinson (2015).

income, wealth & consumption in each case. I assume, that the *ceteris paribus* assumption will hold in the short to medium run and that the approximation is accurate close to the calibrated stationary model.[40]

The values reported in the tables are elasticities, calculated as the percentage change of the level of the variable of interest, relative to its baseline value, divided by a percentage change in risk.

Labour Income								
	Mean				Inequality			
	σ_{Cog}	σ_{Phy}	λ_{Cog}	λ_{Phy}	σ_{Cog}	σ_{Phy}	λ_{Cog}	λ_{Phy}
Overall	0.22	0.07	0.68	0.39	0.71	0.14	0.25	-0.11
Managers	0.23	0.07	0.71	0.36	0.86	0.13	0.25	-0.11
Admin	0.25	0.05	0.79	0.26	0.92	0.05	0.20	-0.14
Trades	0.20	0.08	0.58	0.51	0.68	0.33	0.29	0.05

Table 9: Risk Elasticities - Labour Income

Consumption								
	Mean				Inequality			
	σ_{Cog}	σ_{Phy}	λ_{Cog}	λ_{Phy}	σ_{Cog}	σ_{Phy}	λ_{Cog}	λ_{Phy}
Overall	0.16	0.05	0.66	0.40	0.58	0.12	0.22	-0.10
Managers	0.16	0.05	0.68	0.38	0.76	0.14	0.20	-0.05
Admin	0.14	0.05	0.71	0.34	0.59	0.04	0.04	0.00
Trades	0.15	0.06	0.59	0.47	0.48	0.19	0.25	-0.06

Table 10: Risk Elasticities - Consumption

An analysis of these elasticities, suggests that overall the cognitive skill component of risk (σ_{Cog} & λ^{Cog}) tends to have a bigger impact on aggregate outcomes than corresponding changes in the domain of physical skills. Furthermore changes in technology (λ^{Cog}& λ^{Phy}), tend to have a larger impact on the mean of assets or consumption than on the corresponding Gini coefficient. Whilst the opposite is true for changes in skill risk (σ_{Cog} & σ_{Phy}).

Overall, the mean-elasticities of all three variables with respect to all sources of risk appear to be similar in magnitude, however, when it comes to inequality, wealth appears to be much less responsive than both income and consumption. This suggests that none

[40] i.e. a small change in $\lambda^{Cog}/\lambda^{Phy}$ will not cause a major change in the allocation of workers across occupations, and a small change in $\sigma_{Cog}/\sigma_{Phy}$ will not drastically change the mean or persistence of the skill process.

Wealth								
	Mean				Inequality			
	σ_{Cog}	σ_{Phy}	λ_{Cog}	λ_{Phy}	σ_{Cog}	σ_{Phy}	λ_{Cog}	λ_{Phy}
Overall	0.21	0.05	0.69	0.36	0.16	-0.05	0.03	-0.07
Managers	0.20	0.03	0.71	0.32	0.22	-0.04	0.03	-0.05
Admin	0.30	0.03	0.77	0.26	0.08	-0.04	-0.03	-0.01
Trades	0.21	0.13	0.61	0.53	0.07	-0.08	0.04	-0.11

Table 11: Risk Elasticities - Wealth

of the suggested policy channels will be very effective in addressing wealth inequality, even though they might be effective against income inequality.

The decomposition allows for some interesting thought experiments. For example, by enacting legislation, to reduce σ_{Cog} (e.g. a schooling reform), a government could effectively inequality along all dimensions, but at the same time, the model predicts that the very same policy will cause a reduction in overall savings and therefore hurt production, incomes and ultimately consumption. Uncovering unintended consequences such as these might be one of the potential applications of this model.

Interestingly, encouraging technologies that make productive use of physical skills, has the potential of reducing inequality overall, whilst also increasing income, savings & consumption.

This section has focussed on assessing the impact of different sources of risk on income, wealth and consumption. In the next section, I will focus on the effect of technological change on the economy.

9 Cognitive Skill-Biased Technological Change

9.1 Heterogenous Technological Change

The following presents a simple application that exemplifies the usefulness of the model to study the impact of technological change on the income and wealth distribution. In particular, I investigate the effects of cognitive skill-biased technological change over the period 1980 - 2016 in the UK, using the model and see if it can capture the trends that we observe in the data.

A growing number of papers has documented changes in the demand for (and the return to) different work tasks.[41] The general consensus falls on the side of rising demand

[41] See for example Autor et al. (2003), Autor & Price (2013), Firpo et al. (2011), Rohrbach-Schmidt

for non-routine/complex/analytical/cognitive skills for the period of the 1980s-2000s (c.f. Autor et al. (2003), Spitz-Oener (2006)) with potential tampering off since (c.f. Beaudry et al. (2016)). Conversely, the demand for routine/manual tasks has been falling throughout this period. These general trends are usually attributed to the rising importance of information & communication technologies which complement the former and displace the latter. These changes can have substantial effects on income inequality. For example, in a recent paper, Burstein et al. (2019) argue that computerization can account for as much as 80% of the rise in the skill premium in the US between 1983 and 2003.

Although the computerization of the workplace began in the 1980s and the internet started to permeate into the wider economy in the early 2000s, it is clear to many that the information age is merely in its infancy. New technologies, such as artificial intelligence, big data and machine learning are set to radically transform all sectors of the economy. Given the economic insecurity and potential social upheaval, associated with the move to the *Economy 4.0,* it is of great importance to model and study the impact of technological change on the distribution of income and wealth.

Since Krussell & Smith (1998) technological change has also been an important research area in the heterogenous agent literature (c.f. Heathcote et al. (2009) for a survey). This literature has traditionally studied the effect of aggregate uncertainty created by shocks to total factor productivity. Adding aggregate productivity shocks, however, does little to help us understand the potential effects from task-biased technological change. This is because these affect all workers uniformly, whilst the central finding associated with the research of David Autor and others is that technological change might affect different workers differently. The canonical model of Skill Biased Technological Change (SBTC) for example features two terms of labour augmenting technology: one augmenting the productivity of high-skilled workers and one augmenting the productivity of their low-skilled counterparts.[42]

Labour productivity in this model follows a similar structure. In particular, we can think of the labour productivity in the model, being summarized by the set of skill weights:

& Tiemann (2013), Antonczyk et al. (2009), Bisello (2013), Spitz-Oener (2006)).

[42] See Acemoglu & Autor (2011) for an extensive exposition.

$$\lambda = \underbrace{\begin{bmatrix} \lambda_1^1 & \dots & \lambda_1^N \\ \vdots & \ddots & \vdots \\ \lambda_M^1 & \dots & \lambda_M^N \end{bmatrix}}_{Occupations} \left.\vphantom{\begin{bmatrix} \lambda_1^1 \\ \vdots \\ \lambda_M^1 \end{bmatrix}}\right\} Skills$$

Here the rows summarize different skill groups or occupations and columns represent different work tasks.[43] Since technology only affects labour productivity in this model,[44] the aggregate state of technology is summarized by λ. Rather than being a single aggregate, however, technology in this model is highly heterogeneous, with labour productivity differing across the two dimensions of occupations and tasks.

This provides us with a flexible framework that can not only accommodate Skill or Task biased technological change, but a continuum of hybrid cases, which will be essential for the study of CBTC since I can selectively affect the productivities of different skills and different occupations and study the resulting distributional consequences.

It should be noted, however, that these analyses do not account for changes in the supply of different skills via changes in the skill process. It is therefore only possible to analyze partial equilibrium effects. However, under rigid education markets, these effects might provide a reasonable approximation to short-run dynamics. In the following exercise, I assume that the parameters of the skill process (1.33a) do not change in response to the changes in technological conditions. Given that we are considering very broad, general skills this assumption might not be very strong. It appears reasonable to assume that the mean or the variance of cognitive ability across the population will not dramatically change over the time period ($\sim$ 35 years). The same is possibly true for physical health.

To solve the dynamic model, I make use of the technique developed by Boppart et al. (2017)[45], which has become popular for solving heterogeneous agent models with aggregate shocks. Unlike other papers, however, I only solve for the deterministic transition path from one equilibrium to the next, since I am interested in investigating permanent technological changes, rather than Business Cycle fluctuations.

[43] Every column represents a set of skill weights $\lambda^n \in \lambda$.
[44] Technological change is Harrod neutral.
[45] See Appendix for a description.

9.2 Calibrating the Dynamic Model

The increased usage of computers and related technologies has probably been the most wide-reaching transformation of the economy since the second world war. Yet, the effects of the ICT revolution on the labour market are far from unambiguous: computers are said to augment some tasks, whilst substituting for others as well as generating new fields of work activity (c.f. Acemoglu et al. (2014), Gallipoli & Makridis (2018)).

The theory of TBTC holds, that in general technological change affects the demand for different work tasks and the associated skills respectively. This implies, that a change in technology, benefitting a certain skill (for example the IT revolution favouring cognitive skills), should affect wages proportional to how important the relevant skill is in a certain occupation (c.f. Autor et al. (2003), (2008), Acemoglu & Autor (2011)). I make use of the hypothesized relationship between Computer/IT usage and cognitive skill productivity, to calibrate a path for the evolution of the cognitive skill weights:

1. I use 4 waves[46] of the UK Skills and Employment Survey to obtain a measure of how intensely computers are used within a given occupation cluster in a given year. For this, I rescale the Likert-Scale answers to the survey question "IMPORTANCE OF: USING A COMPUTER/ PC/ OTHER COMPUTERISED EQUIPMENT" on the interval $[0, 1]$[47], and take the mean across the occupation clusters for each year. In the following, I will refer to these values as pc_t^n i.e. the relative value of computer usage in occupation n in year t. The resulting values can be seen in the table below:

	1997	2001	2006	2012
Managers	0.69	0.81	0.90	0.93
Admin	0.75	0.91	0.89	0.93
Trades	0.38	0.38	0.54	0.52

Table 12: Computer Usage over Time

The values in the table provide an interesting insight into the way that computerization occurred across different occupations. For Managers, computer usage seems to have steadily increased by about 10% every 5-year interval, while Administrative & Sales workers experienced a big 20% jump around the turn of the millennium. Skilled workers, Machine Operators, Services and Misc. Occupations experience an even bigger jump (roughly 30%) between the 2001 and 2006 waves. I obtain values for the years not covered by the survey, by using interpolation between the available data points, using a logistic

[46] 1997, 2001, 2006, & 2012 - these are the only years that include a question on computer usage.

[47] With 0 referring to essentially no computer usage and 1 corresponding to the case where computers are essential in the performance of one's work duties.

functional form, which generates the typical "S" shape of technology adoption, that is familiar from work on technology diffusion (c.f. Griliches (1957)).

The following figure depicts the implied adoption curves.[48]

Figure 2: Adoption of ICT by occupation cluster

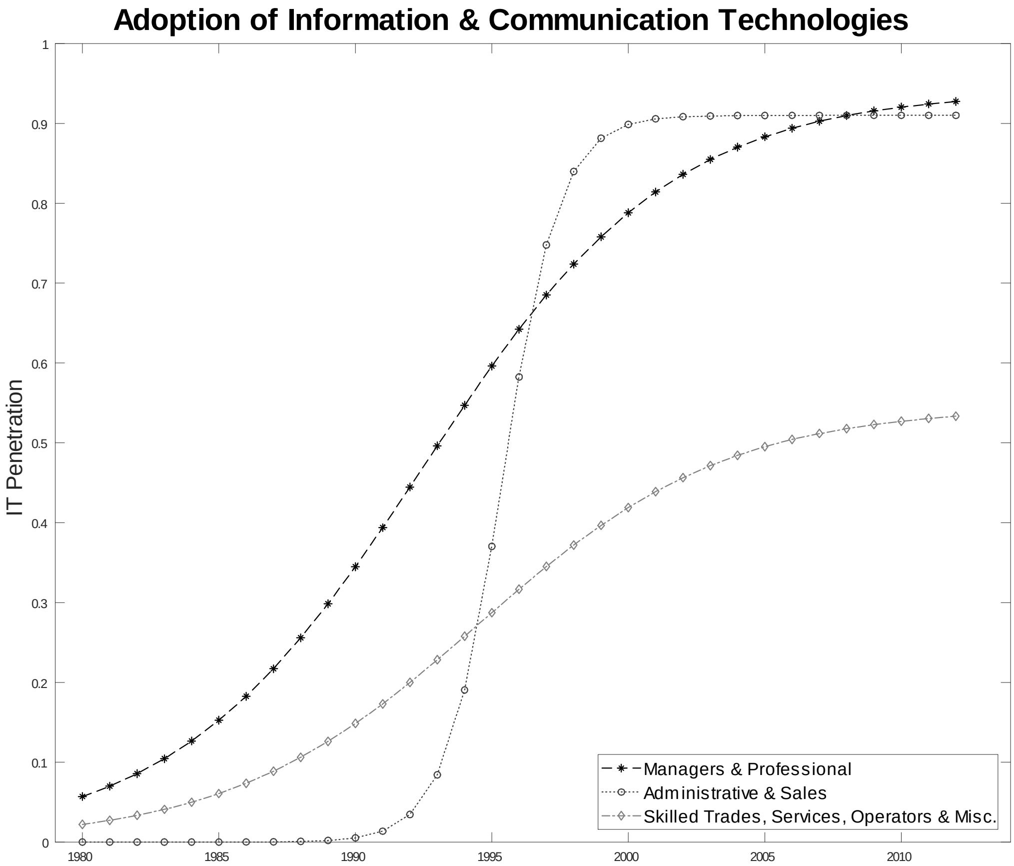

Note: Logistic fit based on available data points. See Table 12.

Source: UK Skills and Employment Survey (1997 - 2012).

2. In order to obtain values for the cognitive skill weights, I assume that computer usage affects the productivity of cognitive skills *log-linearly* so that a percentage change in computer usage is roughly proportional to a $\varkappa^n$ percentage change in cognitive skill productivity:

$$\ln(\lambda^n_{Cog,t}) \approx \varkappa^n pc^n_t \tag{1.42}$$

[48] Obviously there is a lot of uncertainty surrounding, specifically, the earlier years of the sample period for which we have no information from the SES. However, the general patterns seem reasonable; e.g. the advent of computers in the Administrative cluster coincides with the introduction of Microsoft Office in 1990.

Assuming that $\varkappa^n$ is time-invariant, I can eliminate $\varkappa^n$ by dividing both sides of the equation:

$$\frac{\exp^{pc^n_{t+1}}}{\exp^{pc^n_t}} \approx \frac{\lambda^n_{Cog,t+1}}{\lambda^n_{Cog,t}} \tag{1.43}$$

Hence, for any value of pc^n_t, I can obtain an approximation of $\lambda^n_{Cog,t}$ via:

$$\lambda^n_{Cog,t} \approx \frac{\exp^{pc^n_t}}{\exp^{pc^n_{2012}}} * \lambda^n_{Cog,2012} \tag{1.44}$$

I use the *5 Quarter Quarterly Labour Force Survey* (5QLFS) to calibrate the year-to-year occupation cluster transition matrices Π^t. In order to retain comparability, I restrict the sample in the same manner as I have done with the Understanding Society sample. Transitions are measured by movements between the First and Final Interview, in the 5 years centred around the relevant year in order to reduce idiosyncratic measurement error.[49]

I also obtain information about the overall and group-specific distribution of incomes, using the value of usual hourly pay and usual hours worked to infer incomes.

9.3 Evaluating Model Performance - Income Inequality

Having calibrated the model and solved for the dynamic transition path, I plot the model responses against those obtained from the data, to evaluate the model's performance. Unless stated differently, all data points have been smoothed using a 5-year moving average, to reduce the impact of measurement error.

First, we evaluate the model's performance with respect to average labour income. In the demeaned data we see a smooth upward trend for the time between 1993 and 2008, and a small decline afterwards, potentially due to the great recession in the wake of the financial crisis. The model captures the upward trend and some of the later slowdown, however, does not produce as sharp a rise. Still, a regression of the QLFS data on the model-generated path reveals a highly significant, positive association, implying that the model is able to explain about 92% of the variation in mean labour income.

[49] E.g. the Transition Matrix for the year 2000, is estimated by the transitions between 1998 and 1999, 1999 and 2000, 2000 and 2001 and 2001 and 2002.

Figure 3: Mean Labour Income

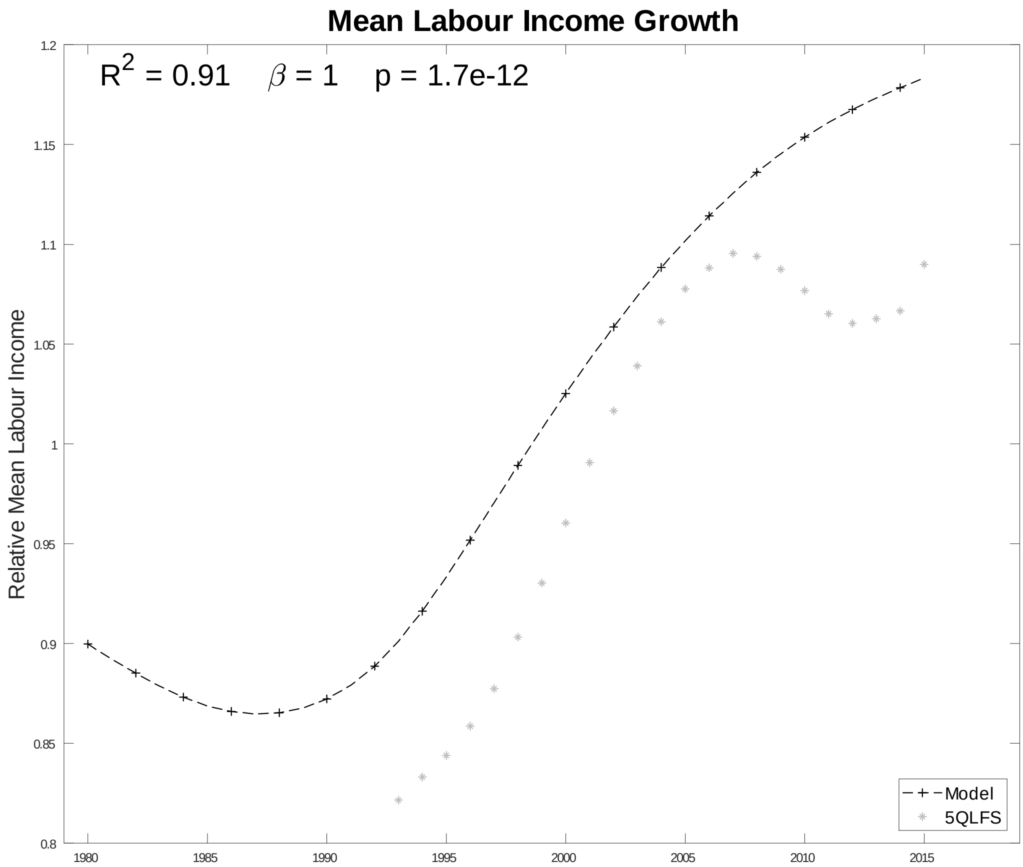

Source: UK QLFS and Model simulations.

Figure 4: Mean Labour Income - Managers & Professionals

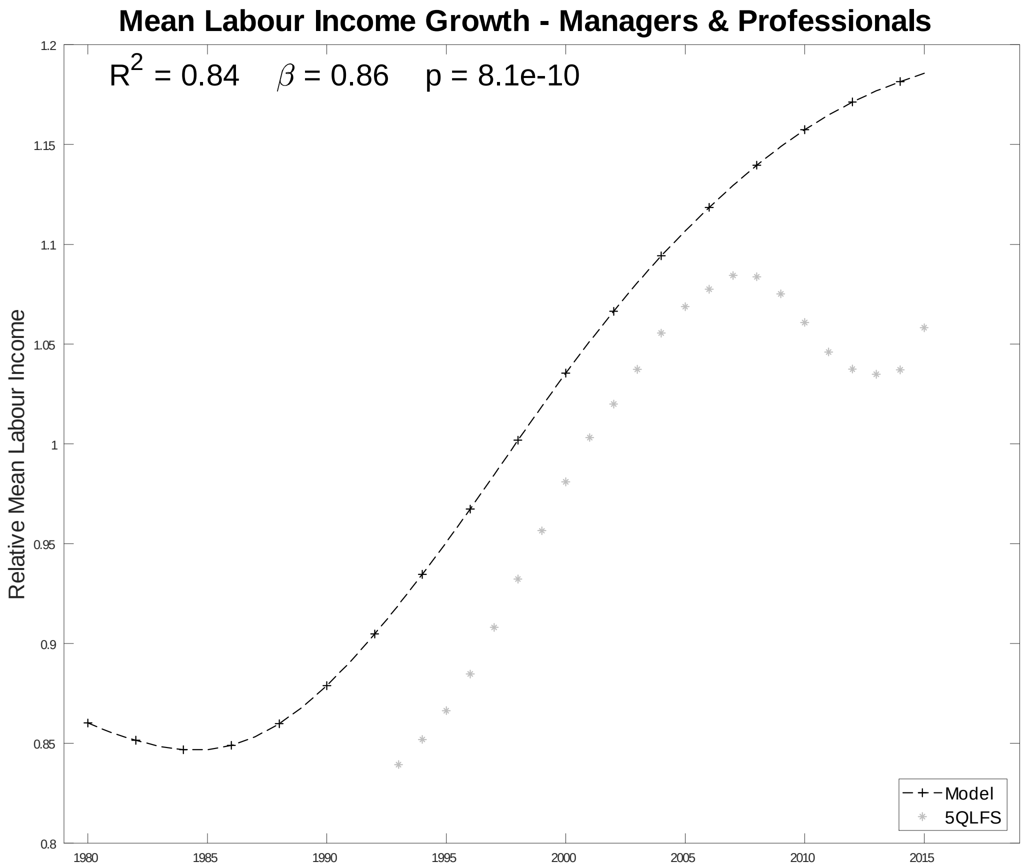

Source: UK QLFS and Model simulations.

Figure 5: Mean Labour Income - Administrative & Sales

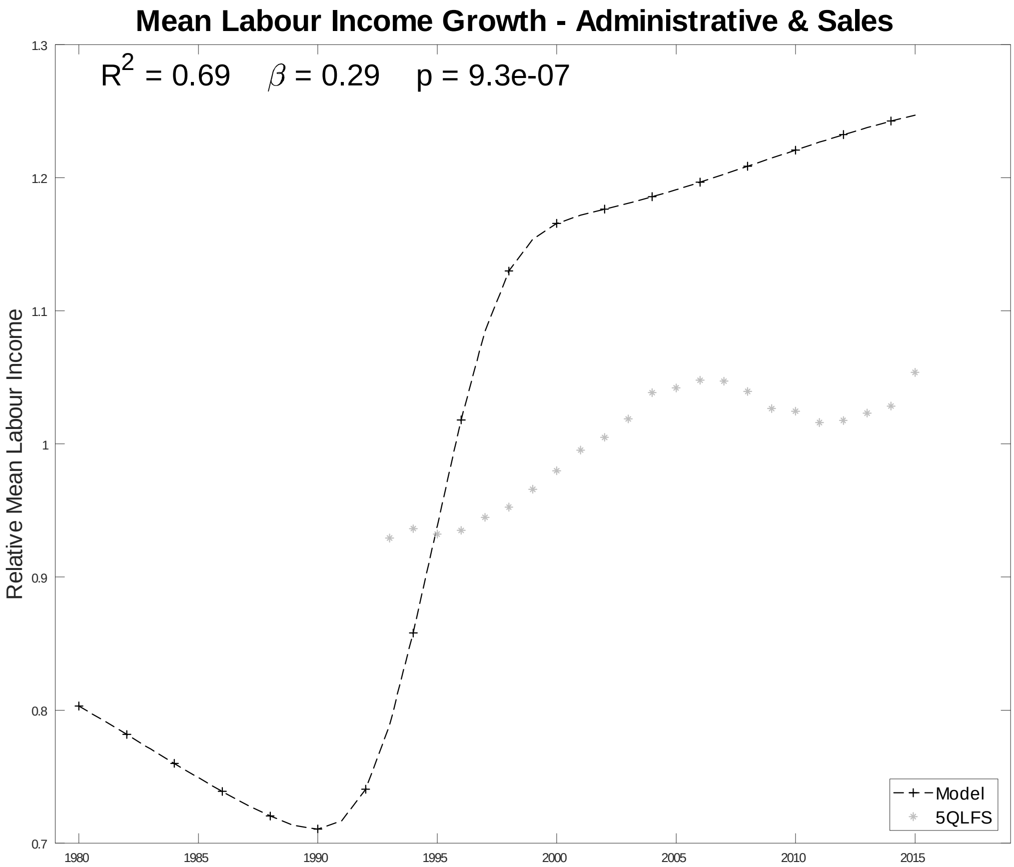

Source: UK QLFS and Model simulations.

Figure 6: Mean Labour Income - Skilled, Services, Operatives & Misc.

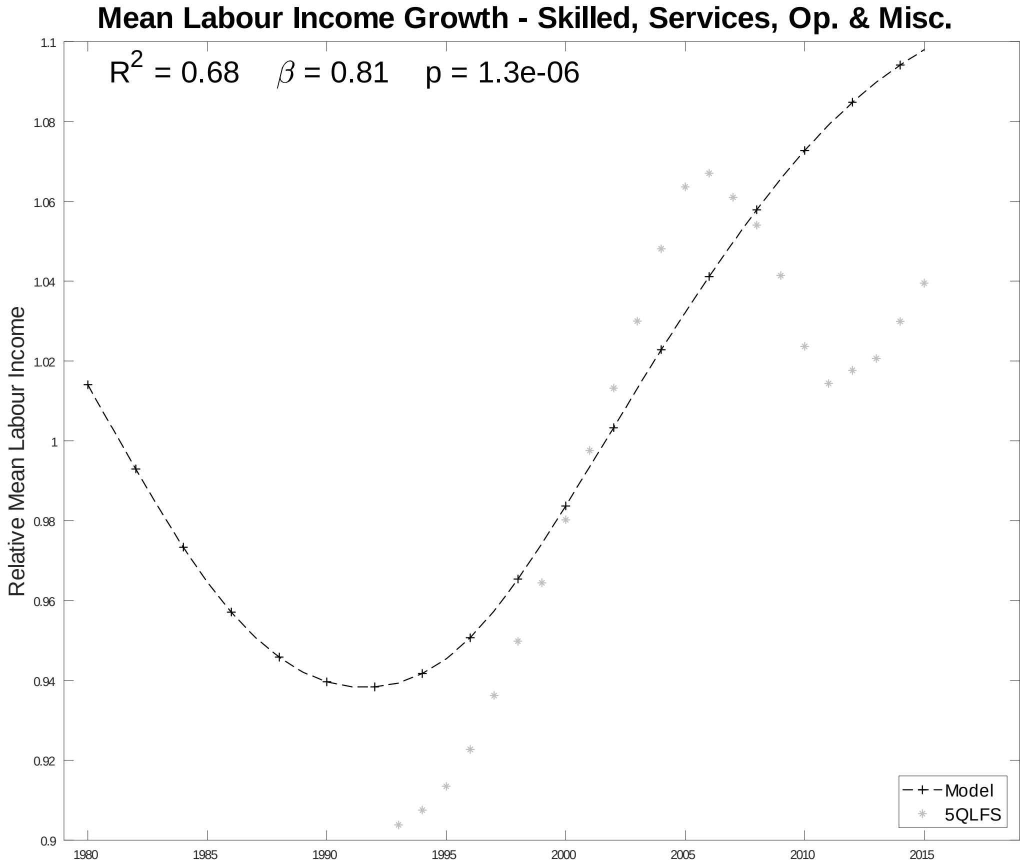

Source: UK QLFS and Model simulations.

The model also does a convincing job in capturing the mean labour income trajectory of Managers, Administrative and Skilled workers, explaining 84, 69 and 68 percent of the variation in the data respectively.

Further, the model generates a clear "Wage Polarization" effect (c.f. Goos et al. (2009), (2014), Goos & Manning (2007)), with the middle-income group of Skilled workers losing their relative position in the income distribution, whilst the lowest paying group (Administrative & Sales workers) catches up over the two decades between 1980 and 2000. Even though the fit with the data is quite weak, the general direction of the model, suggests that a reason for the observed phenomenon of wage polarization could be the comparatively slow computerization of these jobs.

Figure 7: Relative Mean Wages

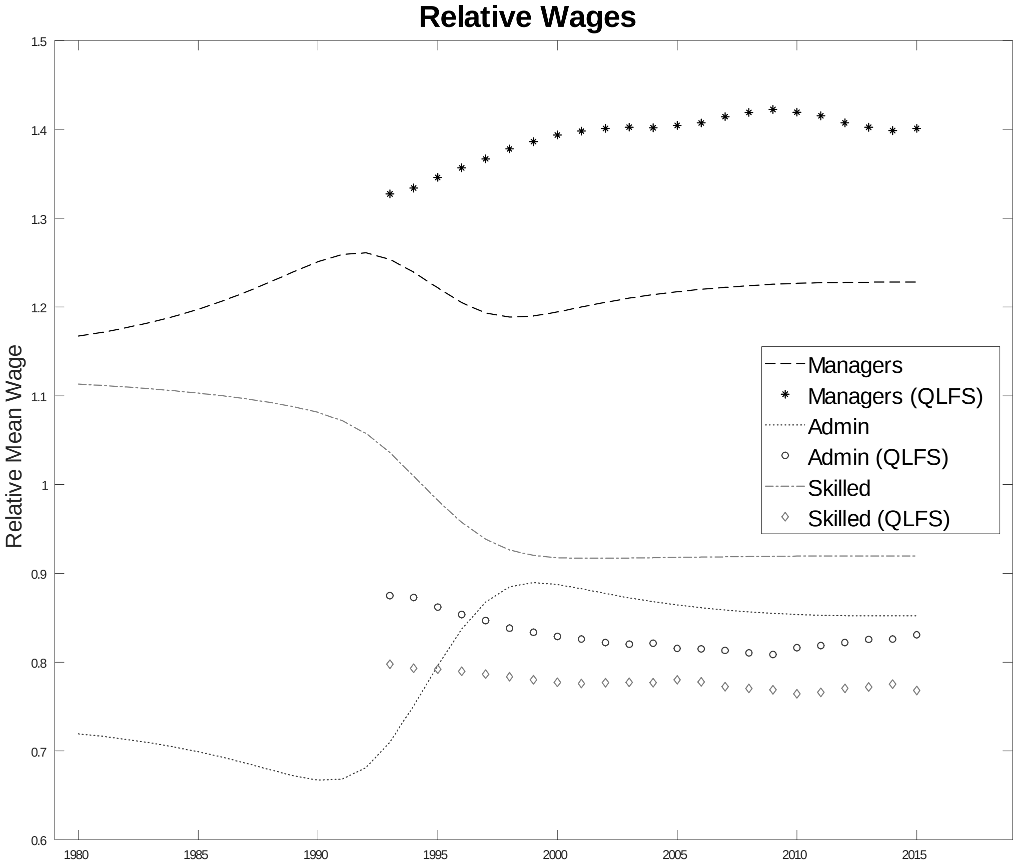

Source: UK QLFS and Model simulations.

The model continues to perform well with respect to income inequality as can be seen in the figures below. The fit is even more apparent, when we consider the relative changes to income inequality, by considering the demeaned values rather than the levels. However, from both, it is apparent that the model manages to account for around 81% of the variation in income inequality as measured by the Gini coefficient. Similar results apply to income inequality as measured by the standard deviation of the logarithms of income. If we consider the possibility, that under sticky wage contracts, changes in labour productivity take some time to show up in measured wages, the fit improves to 88% for a two-year lag.

Figure 8: Labour Income Inequality

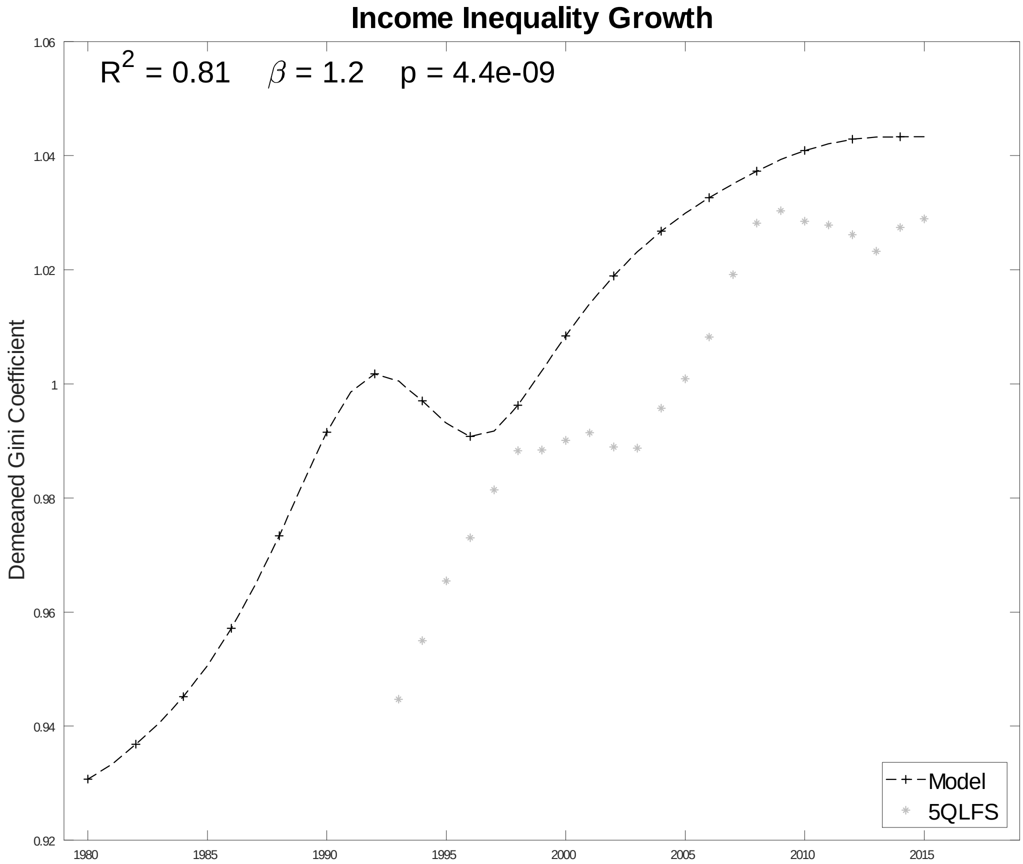

Source: UK QLFS and Model simulations.

Evidently the model's fit is not perfect, and there are many important events, such as the 2007/08 Financial Crisis or the Great Recession that will have affected the degree of income inequality whilst being outside of the model's scope. Generally, however, I hope that this exercise has uncovered some of the general trends due to cognitive skill-biased technological change.

9.4 Evaluating Model Performance - Wealth Inequality

There is generally a lot less information about wealth inequality available, especially for such a specific sample of individuals, so confirming the model's predictions proves much more difficult.

Hills et al. (2013) provide some Gini values for the UK from $1976 - 2005$ and some additional evidence comes from the Wealth & Assets Survey (WAS), started in 2006.

Unfortunately, these two series differ considerably in terms of methodology, so it is difficult to obtain a consistent time series that we can evaluate the model against.

So rather than trying to attempt a dynamic comparison as with income inequality, I will outline the general trends that are reported in these two sources, and see if the model matches the stylized facts.

Hills et al. (2013) report some values on adult, marketable wealth, provided by HMRC for the period $1976-2005$. In terms of the Gini coefficient the pattern that emerges is the following: wealth inequality is stable around 0.65 for the time period $1976-1995$, whereupon it jumps up to 0.71 in 2000 and remains around that value (0.70) in 2005.

The biannual WAS provides information on aggregate total wealth. As reported by the Office for National Statistics, the WAS puts total wealth inequality at 0.61 Gini points for the years $2006-2012$ with an increase to 0.63 in the period $2012-2014$ and a slight drop to 0.62 for the two years after that.

Putting these together we can infer two stylized facts about the path of wealth inequality in the UK:

1. A rapid increase in wealth inequality throughout the 1990s, with a peak somewhere around the year 2000.

2. A plateau stage in the first decade of the new millennium, with a (temporary) small bump in the aftermath of the GFC and the Great Recession.

I will use these two stylized facts as criteria to evaluate the model's performance.

Again, we should not expect the model to perfectly capture the data. Rather I am interested to see whether the forces exerted by **CBTC** are *qualitatively* consistent with the observables.[50]

[50] Generally there are many forces shaping the distribution of wealth and assets that are difficult to capture with standard models. In particular aspects such as the top 1% share of wealth that has become prominent in the popular imagination of wealth inequality is notoriously difficult to reproduce (c.f. De Nardi (2015)).

Figure 9: Wealth Inequality

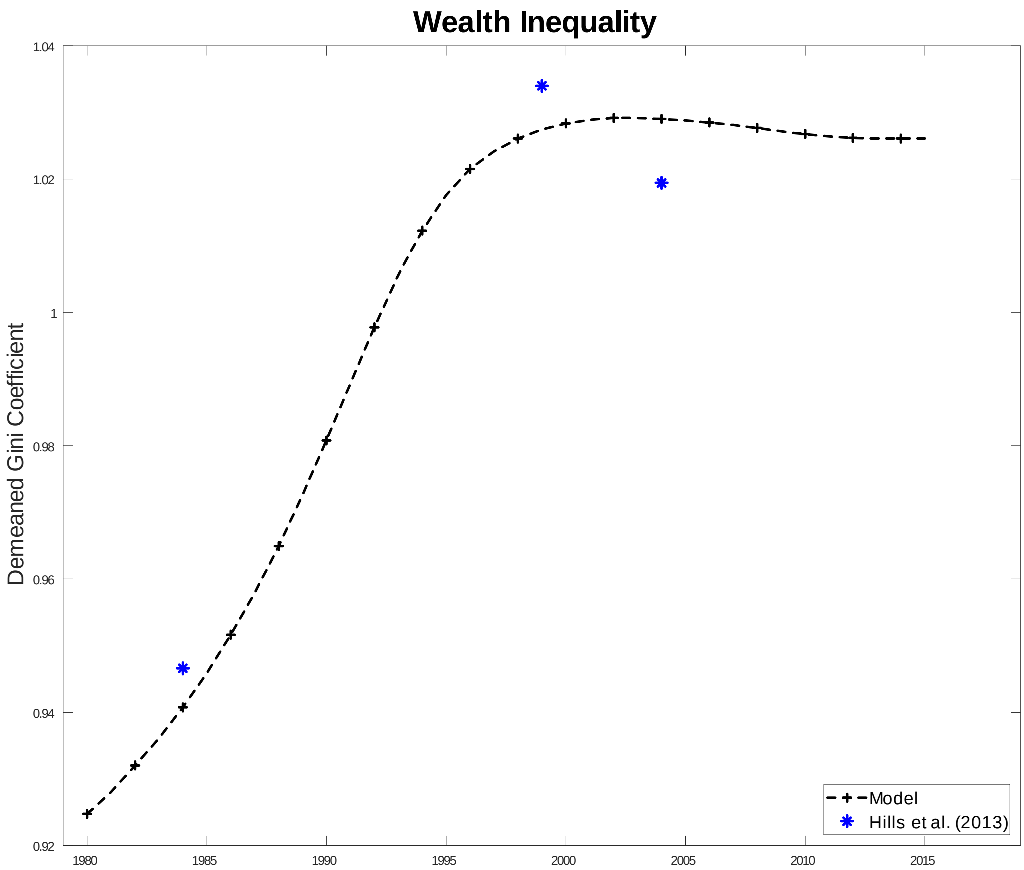

Source: Hills et al. (2013) and Model simulations.

As can be seen from the figure above, the model does a decent job at capturing the two stylized facts. There is a substantial increase in wealth inequality, beginning in 1980, peaking in the year 2000, with a levelling off afterwards.[51] Quantitatively the effects are also quite close. Between 1980 and 2000, inequality increased by 9% in the data, which is matched by an 11% increase predicted by the model. In general, the model translates CBTC into a sizable increase in wealth inequality, which is consistent - if not perfectly - then at least with some of the available evidence (c.f. Roine & Waldenstroem (2015)).

10 Policy Experiments

Having shown that the model is able to capture the general effects of technical change on income and wealth inequality, I will use this section to perform some speculative policy

[51] For obvious reasons the model will not capture the effects of the GFC and the Great Recession.

experiments. These experiments will showcase some potential technological and policy scenarios, in which the model is used to tease out the (partial) effects on income and wealth inequality.

For these purposes I will set a hypothetical time path for λ that corresponds to the changes due to technological progress or policy actions, and perform the same analysis as in the previous section, with the difference, that this time the starting point is the present and I am projecting the economy into the future.

It should be noted that - as with any other prediction exercise - the results of these policy experiments are reliant on a huge number of assumptions and are therefore unlikely to be a *correct* representation of future conditions. Rather than those they should be seen as indications of general trends, that will prevail *ceteris paribus*.

10.1 Reversing Automation

One of the phenomena most closely associated with technological change, was the increasing automation of predominantly routine manual occupations, leading to a considerable fall in the incomes of traditional working-class demographics (c.f. Acemoglu & Autor (2011), Autor et al. (2003), Katz & Murphy (1992), Goldin & Katz (2007)). Apart from the purely economic considerations, automation has wide-ranging social and political implications: Frey at al. (2018) for example show that in the 2016 US election, districts which were exposed to a higher degree of automation were significantly more likely to vote for Donald Trump. Sooner or later it is likely, that the political establishment will respond to the needs of the losers from technological progress, to avoid further disaffection of this large share of the electorate.

For this scenario, I assume that the government implements a set of policies, aimed at artificially increasing the demand for Skilled Trades, Personal Services, Machine Operators and Misc. occupations, either through trade protections, public employment programmes or industry subsidies (c.f. Atkinson (2015)).

In the framework of the model, this is equivalent to a row shift in the λ matrix: the price of every unit of output by these occupations increases due to the government's intervention, effectively increasing the return to each type of skill in these occupations by a common factor. For this scenario, I assume that the government implements a mix of policies, that immediately raises the mean income of the affected group by 20%, whilst all other factors remain unaffected.

The figures below highlight the response of the endogenous variables to this fiscal policy shock.

The immediate effect is a reversal of *Wage Polarization* as the mean income of the affected workers increases. Correspondingly, the Gini coefficient of labour income falls by approximately 1 point.

Higher incomes lead to increased wealth accumulation from the affected workers, even though within-occupation inequality has not changed. There is some evidence of a small effect on the asset accumulation of Managers and Administrative groups, but the effects are small. Overall the policy reduces wealth inequality by about 0.5 Gini points over the first 10 years and around 1 point after 25 years.

Figure 10: Relative Labour Income - Policy Experiment 1

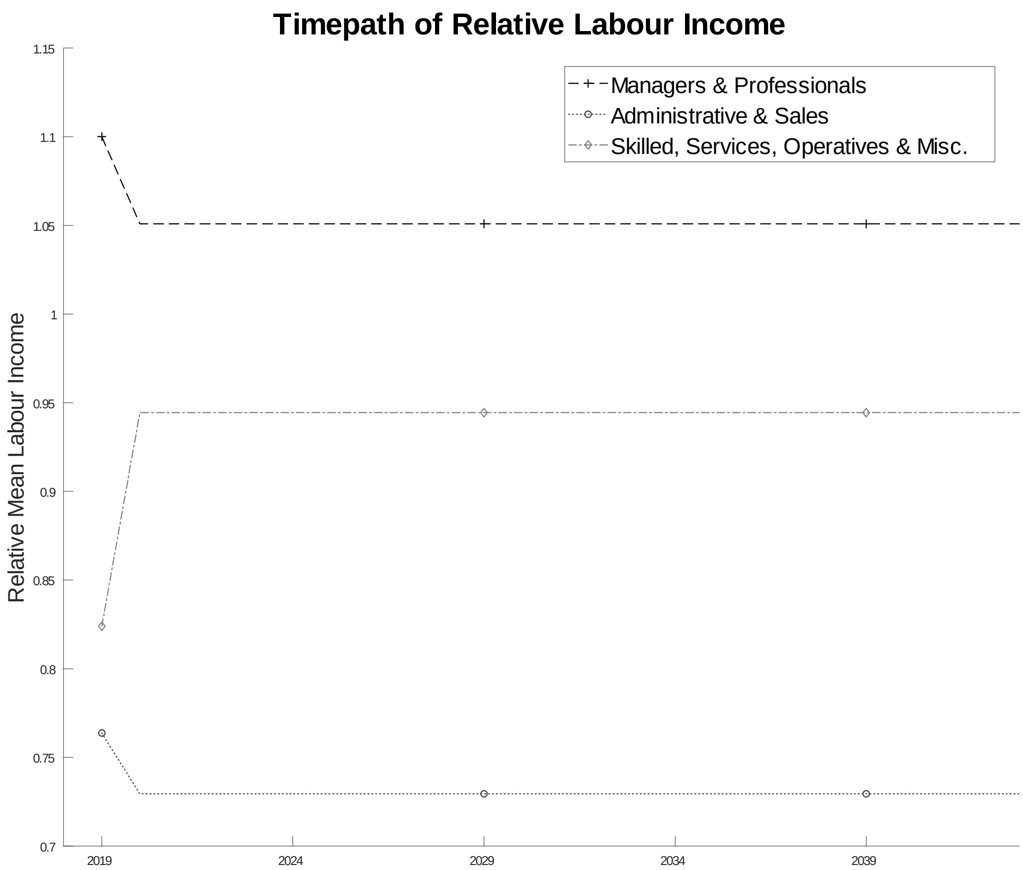

Source: Model simulations.

Figure 11: Occupation Mean Wealth - Policy Experiment 1

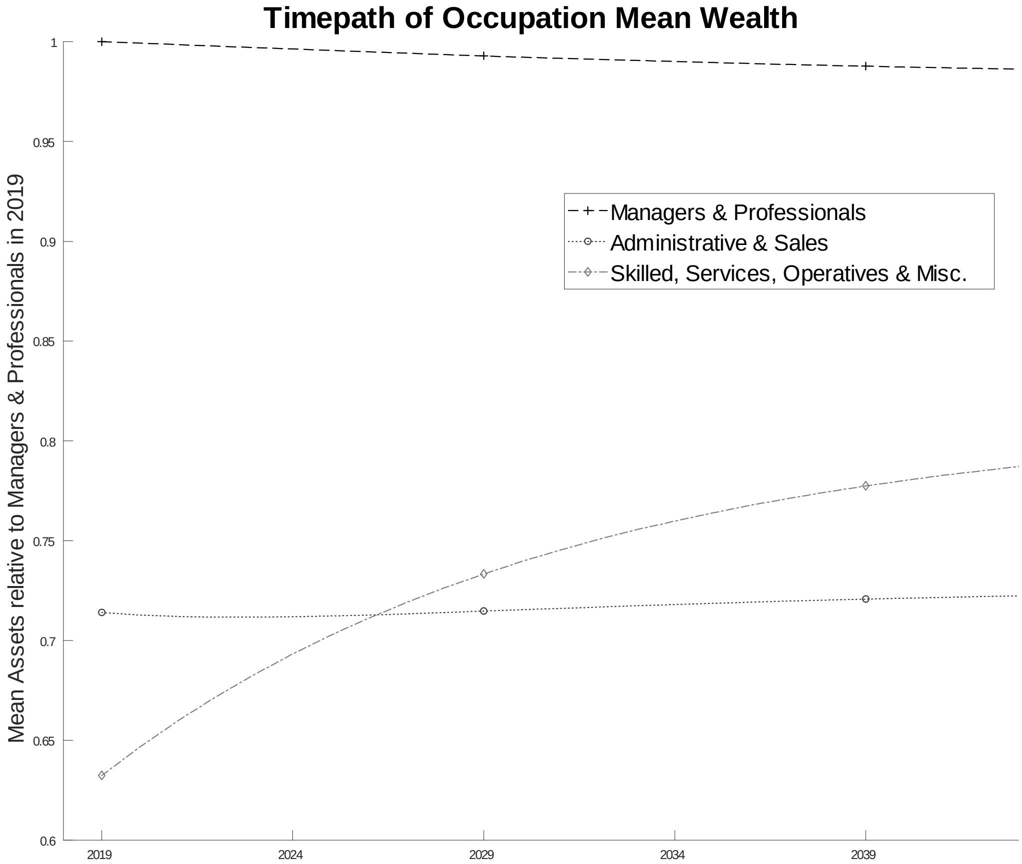

Source: Model simulations.

Figure 12: Wealth Inequality - Policy Experiment 1

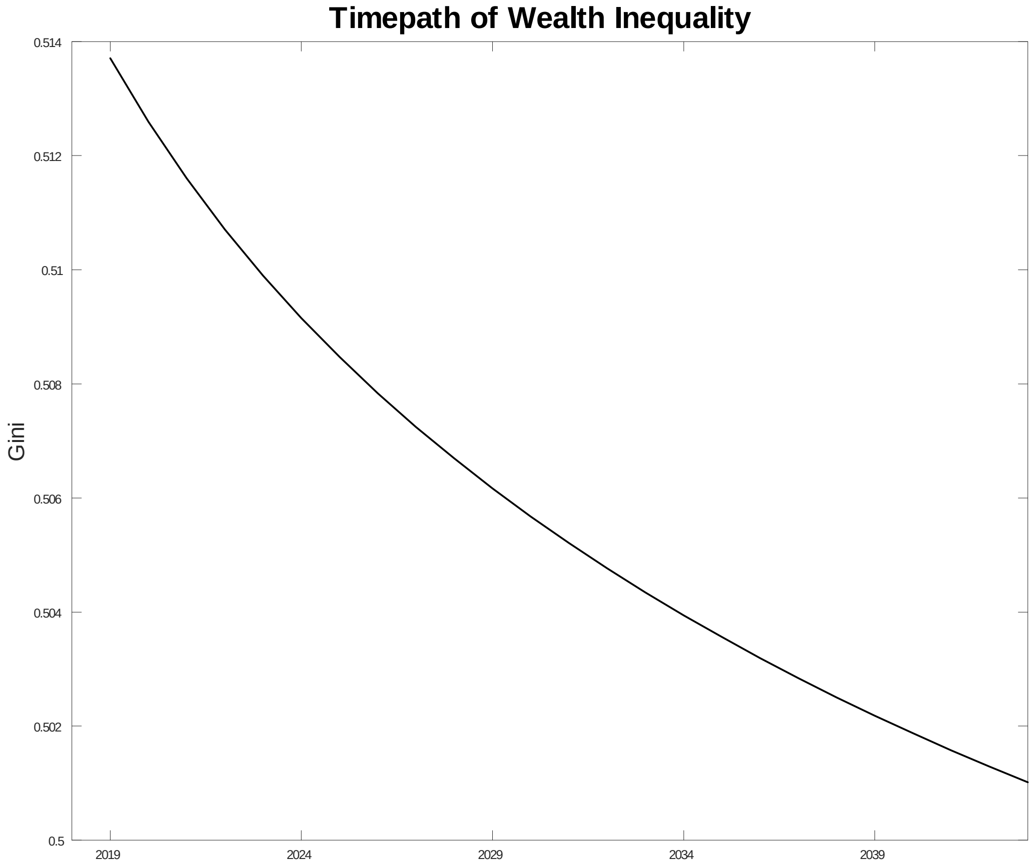

Source: Model simulations.

10.2 Rise of the Machine Brains

With the recent progress in the fields of Machine Learning and Artificial Intelligence, computers are for the first time competing with and outperforming humans in areas such as pattern recognition, logical thinking and strategic decision-making (c.f. Adams (2018)). Traditionally these areas have been seen as bastions of the human brain's comparative advantage, but as Acemoglu & Restrepo (2017) point out, further progress in these areas might see a fall in the wages of those workers we generally see as the most highly skilled.

In this scenario, I simulate the arrival of highly capable artificial intelligence technology, by gradually reducing the cognitive skill weight of Managers & Professionals by 20%.[52] I do not change the return to physical skills, based on the intuition, that the return

[52] I use the following adjustment: $\lambda_{t+1} = (1 - \psi)(\lambda_T - \lambda_t) + \lambda_t$, where $\psi = 0.75$ is an adjustment parameter.

to those tasks that require a physical, human presence (personal interaction, consultations & negotiations, etc.) will not be adversely affected by these developments.

The figures below show the economy's adjustment to technological change. The model indicates a gradual compression of relative wages, as incomes at the top of the distribution contract. Correspondingly there is a gradual fall in overall wage inequality on the order of 0.5 (1) Gini points after 5 (10) years. Asset holdings fall gradually, leading to a slight decrease in wealth inequality (0.5 Gini points over 25 years).

Figure 13: Relative Labour Income - Policy Experiment 2

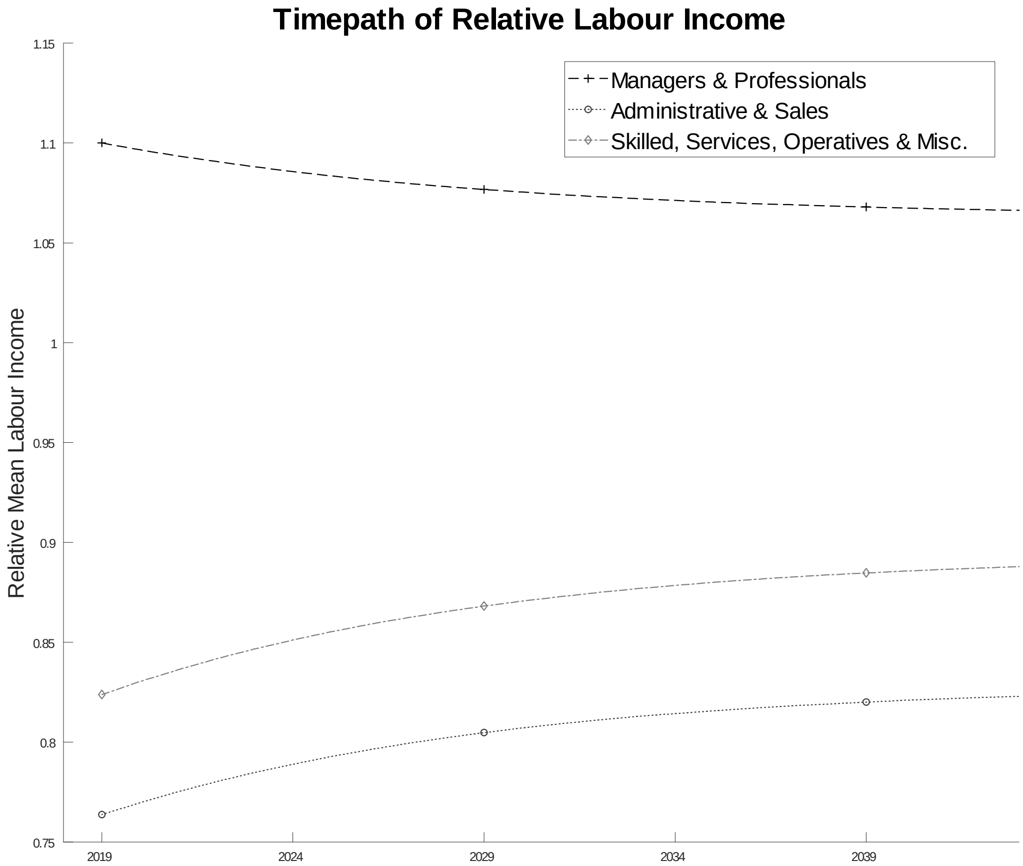

Source: Model simulations.

10.3 The growing importance of "human" skills

The last policy experiment presents itself as the flip side of the previous one. Autor (2015) points out that advances in technology, usually complement those tasks that they don't actively replace. For example, advanced diagnostics technology would free up a doctor's time to consult with patients more thoroughly. Many authors (e.g. Goos (2018)) point out that it is likely that a set of these tasks will be constituted by essentially "human" activities such as social interaction, communication and care. Despite their ability to analyze data rapidly and accurately, a computer's social capabilities are still significantly under-developed.[53] Furthermore, there might be a strong, qualitative preference for interactions with other humans even if the same service was available from a machine. This trend points towards increasing demand for the kind of skills that are commonly called *non-cognitive* or *interpersonal* (c.f. Edin et al. (2022)).

For this scenario I will raise the physical skill weight[54][55] for all occupations where human interaction is likely to be part of the job.[56]

The results of this experiment are fairly similar to the other two experiments: Income inequality is reduced by around 1 Gini point over the first couple of years, as the technological change takes effect. This is accompanied by a slight increase in wealth for the affected groups.

The response of wealth inequality is somewhat interesting, as it exhibits a hump shape - rising initially, and then falling - however, the overall effect is too small to be of much interest.

[53] Showcased by the point that no program to date has managed to pass the Turing test convincingly.

[54] It would have been preferable to have an actual measure of social ability such as the Big5 psychometric measures for this analysis. However in order to stay consistent with the rest of the analysis I simply pick physical skills.

[55] Again, the size of the increase is set to 20% with the adjustment parameter $\psi = 0.75$.

[56] These are Managers & Professionals and Administrative & Sales.

Figure 14: Occupation Mean Wealth - Policy Experiment 3

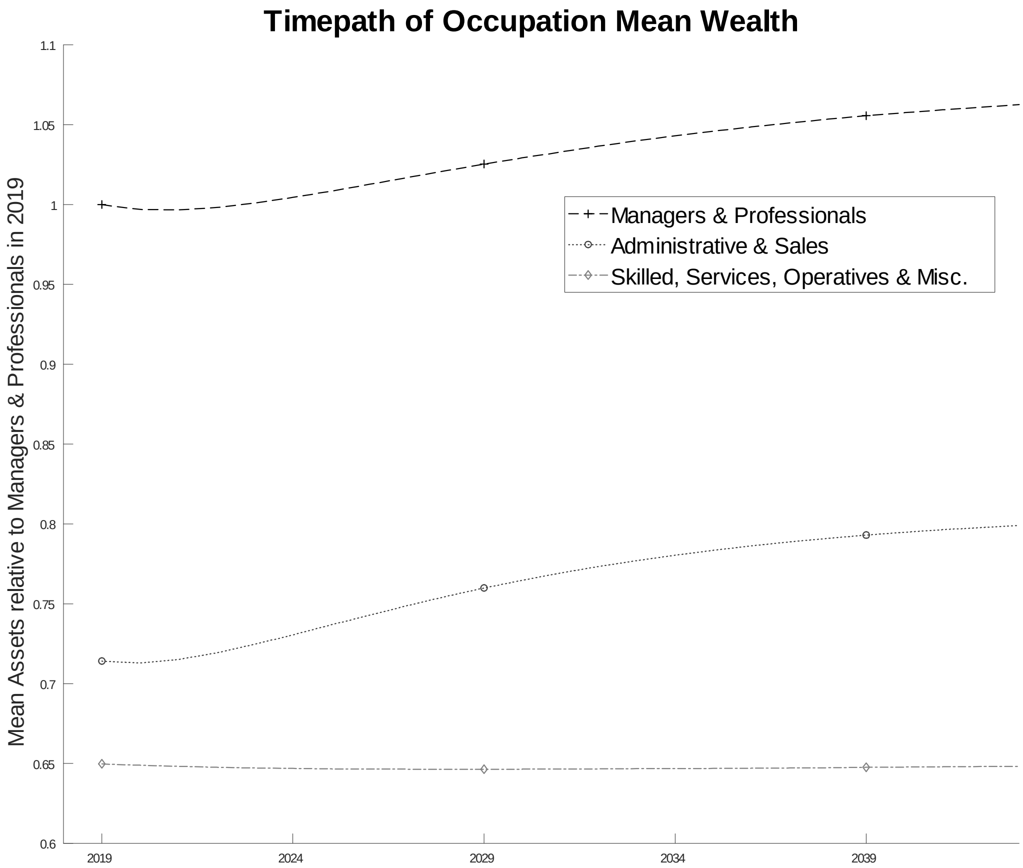

Source: Model simulations.

Figure 15: Wealth Inequality - Policy Experiment 3

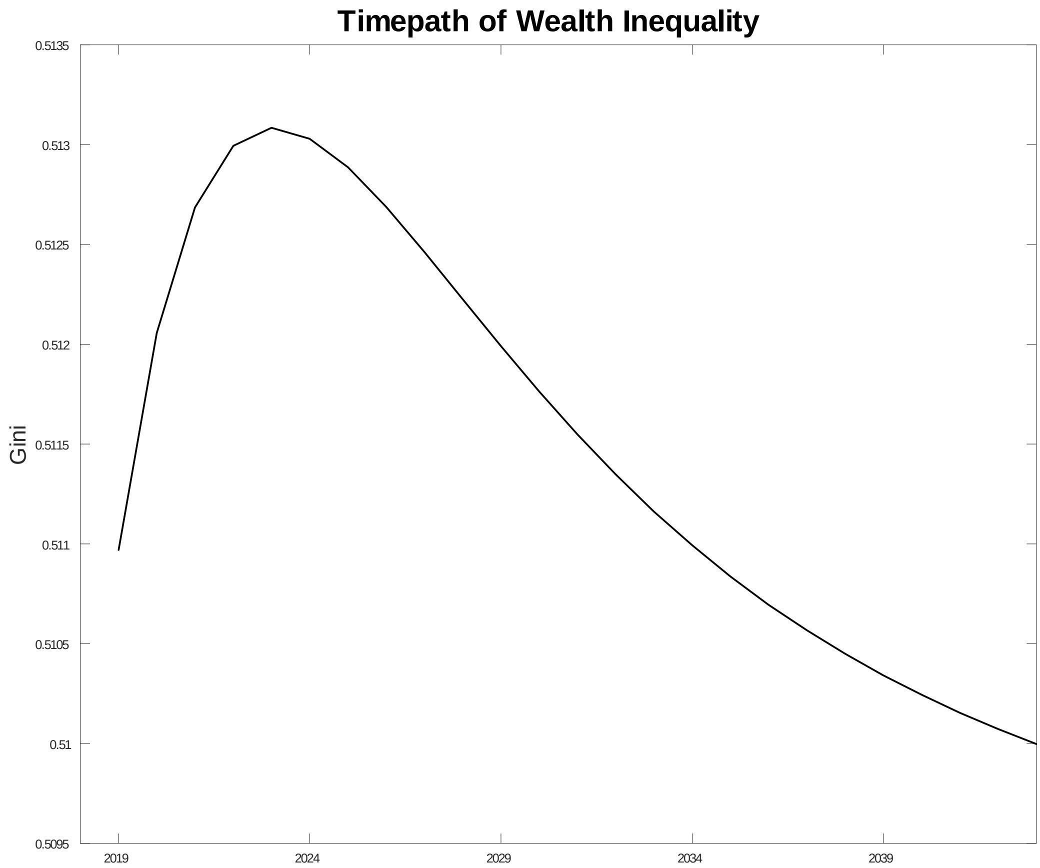

Source: Model simulations.

11 Concluding Remarks

In this chapter, I have extended the standard Aiyagari (1994), model with a structural labour income process, that draws on insights from the wider labour literature. Changes in income are not caused by log-normal productivity shocks, but rather due to combinations of occupational transitions and changes in an individual's skill set. I have shown how to identify the structural parameters from the data showing that the model can reproduce key moments of the empirical income distribution without actively targeting them.

The model has then been evaluated against available data, using a hypothesized relationship between the importance of computer and IT equipment and the productivity of cognitive skills. The results of the dynamic analysis suggest that the calibrated income process is well suited to capture the developments in average wage income as well as the changes to income inequality resulting from cognitive skill-biased technological change.

The analysis suggests that the increased usage of computers and computerized equipment should have contributed to a significant increase in both income and wealth inequality over the period $1980 - 2016$ in the UK. Although it is not possible here to evaluate the implications for wealth inequality sufficiently, this is consistent with some stylized facts and general analyses of the period (c.f. Roine & Waldenstroem (2015)).

I have also evaluated several counterfactual policy & technological change scenarios, assessing the impact on income & wealth inequality.

Overall this suggests that the model could be useful in evaluating the short to medium-run impacts of further technological change on the twin peaks of inequality. With the rising importance of Artificial Intelligence, Automation and the Digital Economy this could be a valuable tool for economists and policy-makers alike. For now, I leave these explorations for future research.

BIBLIOGRAPHY

[1] Acemoglu, D. and Autor, D. "Skills, tasks and technologies: Implications for employment and earnings," Handbook of labor economics (4), 2011, pp. 1043–1171.

[2] Acemoglu, D., Autor, D., Dorn, D., Hanson, G. H. and Price, B. "Return of the Solow Paradox? IT, Productivity, and Employment in U.S. Manufacturing", National Bureau of Economic Research, NBER Working Paper, 2014.

[3] Acemoglu, D. and Jensen, M. K. "Robust comparative statics in large dynamic economies," Journal of Political Economy (123:3), 2015, pp. 587–640.

[4] Acikgoez, O. "On the existence and uniqueness of stationary equilibrium in Bewley economies with production," Journal of Economic Theory (173), 2018, pp. 18–55.

[5] Adams, A. "Technology and the labour market: the assessment," Oxford Review of Economic Policy (34:3), 2018, pp. 349–361.

[6] Aiyagari, S. R. "Uninsured Idiosyncratic Risk and Aggregate Saving," The Quarterly Journal of Economics (109:3), 1994, pp. 659–684.

[7] Angelopoulos, K., Lazarakis, S. and Malley, J. "Wealth inequality and externalities from ex ante skill heterogeneity", 2017.

[8] Antonczyk, Dirk, Bernd Fitzenberger, and Ute Leuschner. "Can a task-based approach explain the recent changes in the German wage structure?." Jahrbücher für Nationalökonomie und Statistik 229, no. 2-3 (2009): 214-238.

[9] Atkinson, A. B. Inequality, Harvard University Press, 2015.

[10] Autor, D. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation", Journal of Economic Perspectives 29(3), American Economic Association, 2015, pp. 3-30.

[11] Autor, D. "The task approach to labor markets: an overview", Journal for Labour Market Research 46(3), Springer Nature, 2013, pp. 185-199.

[12] Autor, D. H. and Handel, M. J. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages," Journal of Labor Economics (31:S1), 2013, pp. S59–S96.

[13] Autor, D. H., Levy, F. and Murnane, R. J. "The Skill Content of Recent Technological Change: An Empirical Exploration," The Quarterly Journal of Economics (118:4), 2003, pp. 1279–1333.

[14] Price, Brendan, and B. Price. "The changing task composition of the US labor market: An update of Autor, Levy, and Murnane (2003)." unpublished manuscript, (2013).

[15] Baumgarten, D., Geishecker, I. and Görg, H. "Offshoring, tasks, and the skill-wage pattern," European Economic Review (61), 2013, pp. 132–152.

[16] Beaudry, P., Green, D. A. and Sand, B. M. "The great reversal in the demand for skill and cognitive tasks," Journal of Labor Economics (34:S1), 2016, pp. S199–S247.

[17] Belfield, C., Blundell, R., Cribb, J., Hood, A. and Joyce, R. "Two decades of income inequality in Britain: the role of wages, household earnings and redistribution," Economica (84:334), 2017, pp. 157–179.

[18] Bewley, Truman. "Stationary monetary equilibrium with a continuum of independently fluctuating consumers." Contributions to mathematical economics in honor of Gérard Debreu 79 (1986).

[19] Bisello, M. "Job polarization in Britain from a task-based perspective. Evidence from the UK Skills Surveys," Department of Economics and Management, University of Pisa Discussion Paper (:160), 2013.

[20] Boppart, T., Krusell, P. and Mitman, K. "Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative," Journal of Economic Dynamics and Control (89), 2018, pp. 68–92.

[21] Burstein, Ariel, Eduardo Morales, and Jonathan Vogel. "Changes in between-group inequality: computers, occupations, and international trade." American Economic Journal: Macroeconomics 11, no. 2 (2019): 348-400.

[22] Carrillo-Tudela, C., Hobijn, B., She, P. and Visschers, L. "The extent and cyclicality of career changes: Evidence for the UK," European Economic Review (84), 2016, pp. 18–41.

[23] Carroll, C. D. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," Economics letters (91:3), 2006, pp. 312–320.

[24] Cortes, G. M. "Where have the middle-wage workers gone? a study of polarization using panel data," Journal of Labor Economics (34:1), 2016, pp. 63–105.

[25] Deming, David J. "The growing importance of social skills in the labor market." The Quarterly Journal of Economics 132, no. 4 (2017): 1593-1640.

[26] De Nardi, Mariacristina. "Quantitative models of wealth inequality: A survey." (2015).

[27] Edin, Per-Anders, Peter Fredriksson, Martin Nybom, and Björn Öckert. "The rising return to noncognitive skill." American Economic Journal: Applied Economics 14, no. 2 (2022): 78-100.

[28] Firpo, S., Fortin, N. and Lemieux, T. "Occupational tasks and changes in the wage structure," (6), 2011, pp. 28.

[29] Frey, C. B., Berger, T. and Chen, C. "Political machinery: did robots swing the 2016 US presidential election?," Oxford Review of Economic Policy (34:3), 2018, pp. 418–442.

[30] Gallipoli, G. and Makridis, C. A. "Structural transformation and the rise of information technology," Journal of Monetary Economics (97), 2018, pp. 91–110.

[31] Gathmann, C. and Schönberg, U. "How general is human capital? A task-based approach," Journal of Labor Economics (28:1), 2010, pp. 1–49.

[32] Gensowski, Miriam. "Personality, IQ, and lifetime earnings." Labour Economics 51 (2018): 170-183.

[33] Goldin, C. and Katz, L. F. "The race between education and technology: the evolution of US educational wage differentials, 1890 to 2005", Technical report, National Bureau of Economic Research, National Bureau of Economic Research, 2007

[34] Goos, M. "The impact of technological progress on labour markets: policy challenges," Oxford Review of Economic Policy (34:3), 2018, pp. 362–375.

[35] Goos, M. and Manning, A. "Lousy and lovely jobs: The rising polarization of work in Britain," The review of economics and statistics (89:1), 2007, pp. 118–133.

[36] Goos, M., Manning, A. and Salomons, A. "Job polarization in Europe," American economic review (99:2), 2009, pp. 58–63

[37] Goos, M., Manning, A. and Salomons, A. "Explaining job polarization: Routine-biased technological change and offshoring," American Economic Review (104:8), 2014, pp. 2509–26.

[38] Gourieroux, Christian, Monfort Gourieroux, Alain Monfort, and Director Alain Monfort. Simulation-based econometric methods. Oxford university press, 1996.

[39] Griliches, Z. "Hybrid corn: An exploration in the economics of technological change," Econometrica, Journal of the Econometric Society (), 1957, pp. 501–522.

[40] Guvenen, Fatih, Serdar Ozkan, and Jae Song. "The nature of countercyclical income risk." Journal of Political Economy 122, no. 3 (2014): 621-660.

[41] Guvenen, F., Karahan, F., Ozkan, S. and Song, J. "What do data on millions of US workers reveal about life-cycle earnings risk?", SSRN Electronic Journal , Technical report, National Bureau of Economic Research, Elsevier BV, 2015.

[42] Guvenen, F. and Kuruscu, B. "Understanding the evolution of the US wage distribution: A theoretical analysis," Journal of the European Economic Association (10:3), 2012, pp. 482–517.

[43] Guvenen, F., Kuruscu, B., Tanaka, S. and Wiczer, D. "Multidimensional Skill Mismatch," (), 2017.

[44] Heathcote, J., Storesletten, K. and Violante, G. L. "Quantitative Macroeconomics with Heterogeneous Households," Annu. Rev. Econ. (1:1), 2009, pp. 319–354.

[45] Heathcote, J., Storesletten, K. and Violante, G. L. "The macroeconomic implications of rising wage inequality in the United States," Journal of political economy (118:4), 2010, pp. 681–722.

[46] Hills, J., Bastagli, F., Cowell, F., Glennerster, H. and Karagiannaki, E. Wealth in the UK, Oxford University Press, 2013.

[47] Hubmer, Joachim, Per Krusell, and Anthony A. Smith. The historical evolution of the wealth distribution: A quantitative-theoretic investigation. No. w23011. National Bureau of Economic Research, 2017.

[48] Huggett, Mark. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." Journal of economic Dynamics and Control 17, no. 5-6 (1993): 953-969.

[49] Huggett, M., Ventura, G. and Yaron, A. "Sources of Lifetime Inequality," American Economic Review (101:7), 2011, pp. 2923–54.

[50] Imrohoroğlu, A. "Cost of business cycles with indivisibilities and liquidity constraints," Journal of Political economy (97:6), 1989, pp. 1364–1383.

[51] Kambourov, G. and Manovskii, I. "Occupational Mobility and Wage Inequality," The Review of Economic Studies (76:2), 2009, pp. 731–759.

[52] Katz, L. F. and Murphy, K. M. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors", 1992, pp. 35-78.

[53] Lazear, E. P. "Firm-Specific Human Capital: A Skill-Weights Approach," Journal of Political Economy (117:5), 2009, pp. 914–940.

[54] Lise, J. and Postel-Vinay, F. "Multidimensional skills, sorting, and human capital accumulation", 2016.

[55] Ljungqvist, L. and Sargent, T. J. Recursive Macroeconomic Theory, MIT press, 2012.

[56] Meghir, Costas, and Luigi Pistaferri. "Earnings, consumption and life cycle choices." In Handbook of labor economics, vol. 4, pp. 773-854. Elsevier, 2011.

[57] Miao, J. Economic dynamics in discrete time, MIT press, 2014.

[58] Milanovic, B. Global inequality: A new approach for the age of globalization, Harvard University Press, 2016.

[59] Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo. "Uneven growth: automation's impact on income and wealth inequality." Econometrica 90, no. 6 (2022): 2645-2683.

[60] Piketty, T. "Capital in the 21st Century," Cambridge: Harvard Uni (), 2014.

[61] Quadrini, V. and Ríos-Rull, J.-V. "Inequality in macroeconomics", Elsevier, 2015, pp. 1229–1302.

[62] Rohrbach-Schmidt, D. and Tiemann, M. "Changes in workplace tasks in Germany—evaluating skill and task measures," Journal for Labour Market Research (46:3), 2013, pp. 215–237.

[63] Roine, J. and Waldenström, D. "Long-run trends in the distribution of income and wealth"'Handbook of income distribution', Elsevier, 2015, pp. 469–592.

[64] Rouwenhorst, G. K. .. Asset pricing implications of equilibrium business cycle models In: Cooley, T.F. (Ed.), Frontiers of Business Cycle Research, Princeton University Press., Princeton University Press, 1995.

[65] Roy, Andrew Donald. "Some thoughts on the distribution of earnings." Oxford economic papers 3, no. 2 (1951): 135-146.

[66] Roys, N. and Taber, C. "Skills prices, occupations and changes in the wage structure," (), 2016.

[67] Sanders, C. and Taber, C. "Life-Cycle Wage Growth and Heterogeneous Human Capital," Annu. Rev. Econ. (4:1), 2012, pp. 399–425.

[68] Slavík, C. and Yazici, H. "Wage risk and the skill premium," (), 2017.

[69] Spitz-Oener, A. "Technical change, job tasks, and rising educational demands: Looking outside the wage structure," Journal of labor economics (24:2), 2006, pp. 235–270.

[70] Stiglitz, J. "The Price of Inequality", 2013, pp. 52-53.

[71] Stokey, Nancy, Robert Lucas, and C. Edward. "with PRESCOTT." Recursive Methods in Economic Dynamics (1989).

[72] Tauchen, George. "Finite state markov-chain approximations to univariate and vector autoregressions." Economics letters 20, no. 2 (1986): 177-181.

[73] Tinbergen, J. "Substitution of graduate by other labour," Kyklos (27:2), 1974, pp. 217–226.

[74] Tinbergen, J. "Substitution of academically trained by other manpower," Review of World Economics (111:3), 1975, pp. 466–476.

[75] Vivarelli, M. "Innovation, employment and skills in advanced and developing countries: A survey of economic literature," Journal of Economic Issues (48:1), 2014, pp. 123–154.

[76] Ware, J. E., Dewey, J. E. and Kosinski, M. How to Score Version 2 of the SF-36 Health Survey:(standard & Acute Forms);[SF-36v2], Vol. 40, QualityMetric, 2001.

[77] Whitley, E., Deary, I. J., Ritchie, S. J., Batty, G. D., Kumari, M. and Benzeval, M. "Variations in cognitive abilities across the life course: Cross-sectional evidence from Understanding Society: The UK Household Longitudinal Study," Intelligence (59), 2016, pp. 39–50.

[78] Yamaguchi, S. "Tasks and Heterogeneous Human Capital," Journal of Labor Economics (30:1), 2012, pp. 1–53.

12 Appendix A

12.1 Cognitive Skill Measures

Wave 3 of Understanding Society contains an additional module financed by the Economic and Social Research Council (ESRC) with resources from the Large Facilities Capital Fund of the Department for Business, Innovation, and Skills. This module added question stages assessing the cognitive and psychological traits of adult (16+) respondents.

In this chapter, I follow Whitley et al. (2016) by selecting 5 exercises from the cognitive module to create a composite measure of cognitive ability. These exercises relate to: 1. Numeric Ability, 2. A Subtraction Exercise, 3. Completion of a Number Sequence, 4. A word recall exercise & 5. Verbal Fluency.

As is apparent from the types of exercises these exercises test the respondents logical, mathematical as well as lingual intelligence. However, it is likely that all these areas are related to what is commonly referred to as *general intelligence*. To uncover this latent factor I perform a principal component analysis (PCA) on the standardized test results. The PCA suggest that only the first principal component has an Eigenvalue greater than 1, which is commonly seen as the cutoff for significance. I interpret this insofar that there is only a single factor that is relevant in explaining an individual's performance in all 5 exercises. I dub this factor *general intelligence* or *cognitive ability.*

The table below reports the factor loadings of the variables. All loadings are positive, suggesting that doing better in any type of test is associated with a higher level of general intelligence, which is what one might expect. The loadings are also of similar size, even though numeric ability and the number sequence exercises have slightly higher loadings.

12.2 Calibration Procedure

In order to calibrate the process for cognitive skills I loosely follow the Method of Simulated Moments (MSM) set out by Guvenen et al. (2014). However, it is probably more

Variable	Loading	Unexplained Variation
Numeric Ability	0.53	0.41
Subtraction Exercise	0.40	0.67
Number Sequence	0.52	0.43
Word Recall	0.39	0.68
Verbal Fluency	0.38	0.69

Table A1: Factor Loadings of PCA on Cognitive Skills

accurate to talk of an Indirect Inference (II) procedure (c.f. Gourieroux et al. (1996)), as the target moments are obtained by running a regression on simulated data. In this case, the target moments (m_1, m_2) are the persistence (κ), and the standard deviation of incomes σ_y, observed in the data, as described in equation (1.41a). The corresponding simulated moments $(\tilde{m}_1(\theta), \tilde{m}_2(\theta))$ are in turn generated by:

1. Picking a vector of parameters θ. In this case $\theta = (\rho_{Cog}, \sigma_{Cog})$, with all other parameters held fixed.

2. Solving the model given those parameters.

3. Simulating a series of labour market histories.

4. Estimating regression (1.41a), to obtain $(\tilde{m}_1(\theta), \tilde{m}_2(\theta))$.

To implement the procedure, I define $F_n(\theta) = \frac{\tilde{m}_n - m_n}{|m_n|}$, i.e. $F_n(\theta)$ corresponds to the percentage deviation of the simulated n^{th} moment from its target, given the parameter vector θ. I then choose θ to minimize

$$\min_{\theta} F'(\theta) W F(\theta) \tag{1.45}$$

where W is the identity matrix, as $F(\theta)$ is already scaled. To solve (1.45) I employ a nonlinear, derivative-free, minimization routine.[57]

12.3 Boppart et al. (2017) Algorithm

The solution method follows Boppart et al. (2017):

1. Choose a time T at which the economy has presumably reached a steady state.[58]

[57] I also double-check, using a genetic algorithm, that the local solution coincides with the global minimum.

[58] For my application I set T= 2200, but most variables settle down after around 20 years.

2. Solve for the stationary equilibrium at T where the skill weights are given by λ_T.

3. Construct the deterministic transition path of $\{\lambda_t\}_0^T$ and $\{\Pi_t\}_0^T$.

4. Make an initial guess for the transition path of capital $\{K\}_0^T$.

5. Given the guess for the transition path, and the evolution of λ, solve the policy functions backwards from $t = T - 1$, setting $G^T = G^{SS}$.

6. Calculate the transition matrix for the joint state Λ in every time period, using the policy functions obtained in step **5.** Iterate the joint distribution forward, starting with the initial stationary distribution.

7. Calculate the implied capital stock at each point in time and update the guess for the path of capital until convergence.

12.4 Occupation Specific Productivities

In this chapter, I sometimes suggested that Managers & Professionals might appear to have high returns to physical skills due to some higher level of unobserved productivity. In the language of the model, this corresponds to a third - general - skill, that every worker is evenly endowed with. For simplicity, I set the skill endowment of this general skill to 1. The skill weight, associated with this general skill, can be interpreted as the "base productivity" of the relevant occupation. Ignoring this base productivity, can in theory bias the estimates of the skill weights.

In this appendix section, I test this hypothesis, by including Occupation Dummies in the skill weights regression. These dummy variables will capture any higher level of productivity that is unrelated to differences in the skill productivities.

The extended econometric model is given by:

$$\omega_{i,n} = \sum_{n=1}^{N} D_n + \sum_{n=1}^{N} \left\{ D_n \sum_{m=1}^{M} \tilde{\lambda}_m^n s_{i,m} \right\} \tag{1.46}$$

The regression suggests that excluding occupation-specific productivities has indeed affected the skill weight estimates. After including the occupation-specific intercepts the estimated returns to all skills have decreased significantly. This is particularly true for Managers & Professionals, who have the highest occupation-specific productivity. The general patterns, however, are largely unaffected: Managers & Professionals still have the highest returns to cognitive skills, even though they are now very close to the estimate for Admin & Sales occupations. Skilled Workers and other physically demanding occupations

	Labour Income
Managers & Professionals	934.429***
	(202.988)
Admin & Sales	-530.172
	(319.563)
Skilled Workers, Services, Operatives & Misc	-95.705
	(234.703)
Managers & Professionals X Cognitive Ability	1523.511***
	(188.932)
Admin & Sales X Cognitive Ability	1393.683***
	(268.147)
Skilled Workers, Services, Operatives & Misc X Cognitive Ability	786.205***
	(121.178)
Managers & Professionals X Physical Ability	501.617*
	(220.936)
Admin & Sales X Physical Ability	314.609
	(291.348)
Skilled Workers, Services, Operatives & Misc X Physical Ability	460.493***
	(133.514)
R-squared	0.212
N	4393

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A2: Skill Weights Regression with Occupation Specific productivities still have the highest return to physical skills, even though the return has dropped by around half. Curiously the physical skill weight of Managers & Professionals is still very close to that of Skilled Workers, suggesting that the previous estimates weren't entirely due to higher unobserved productivity.

12.5 Comparison of Skill Weights with SES-derived measures

The presented approach to estimating the skill weights allows for an interesting comparison with the literature. The common approach to recovering skill weights, is based on survey information on individual job content (c.f. Autor & Handel (2013) or Spitz-Oener (2006)).

For this approach, information on the perceived importance of different tasks is transformed into a measure of the importance of a category of work tasks (such as analytic or manual tasks), and then used as a basis for analysis (see Autor (2013) for an overview of the method and some criticisms).

Since the approach I presented above, differs fundamentally from the commonly employed method it might be useful to cross-validate the estimated skill weights with measures derived from job-task data.

For this purpose, I employ a wave of the UK Skills & Employment Survey. The UK Skill & Employment Survey (SES) is the UK's primary (and only) source of job task information and contains individual-level survey information on different aspects of job content, educational attainment and other demographic factors. I follow Bisello (2013) to select a number of survey items corresponding to analytic and manual tasks and perform a Principal Component Analysis, obtaining the first principal component in each case. The retained principal components are standardized on $[0, 1]$.

For this, I limit the sample to the year 2012, which is closest to the time of the 3rd UnSoc. wave (2011 - 2013). I then compare the skill weights obtained from the regression procedure with those implied by the SES data.[59] In an ideal scenario, where workers have a perfect grasp of the contribution that the different work tasks they perform make towards their overall productivity, we should expect a high level of correlation between the skill weights derived from the SES and those estimated from the UnSoc data.

The correlation between the measures for cognitive and analytic skill weights is 0.66, which warrants some confidence in our measures. The correlation for physical and manual skill weights, however, is 0.04 which, at first glance, appears puzzling. A couple of explanations suggest themselves:

[59] In order to increase the number of observations I use the skill weights obtained for all 9 major SOC2000 occupations.

1. The SF-12 PCS measure does not correspond well with the tasks covered by the manual skills measure, obtained from the SES.

2. Individuals are unable to properly assess the importance of their manual work activities, or the self-reported PCS measures are biased in some unknown way.

3. The correlation might be driven by high estimated physical skill weights for some occupations (e.g. Managers or Professionals) which might in part reflect higher overall productivity in these occupations, rather than a strictly higher importance of physical skills.

To assess this possibility I take the relative importance of skills, by forming the ratio of cognitive/analytical and physical/manual skill weights. The correlation between the transformed variables is 0.35, which still leaves a lot unexplained, suggesting that one of the other explanations is contributing to the mismatch.

Overall this little comparison has provided us with some new and interesting insights: It seems that survey-based approaches recover similar skill weights as regression-based approaches, at least for cognitive or analytic tasks. This is an important insight since it suggests that at least in this area perceived importance somewhat accurately reflects true labour productivity.

CHAPTER 2

THE INTERPLAY BETWEEN WEALTH AND HUMAN CAPITAL INEQUALITY – IMPLICATIONS FOR THE UK'S POST-COVID-19 RECOVERY

1 Introduction

Wealth and human capital are closely intertwined (Lazear (1975)). Accumulated wealth allows individuals to invest in costly education and training, which in turn increases their potential labour earnings. It also provides a safety net in the event of unforeseen circumstances or economic shocks, allowing individuals to maintain their human capital investments. Additionally, wealth provides the opportunity to take risks and invest in skills that may not provide immediate returns but can pay off in the long run. In turn, human capital aids in the accumulation of wealth, by enabling more skilled workers to earn higher wages in the labour market, or use their financial sophistication to make better investment choices and grow their savings.

This relationship between wealth and human capital has far-reaching implications for economic development, inequality and social mobility. Wealth disparities can perpetuate and amplify existing inequalities, as those with higher levels of wealth can access more expensive education, training and other opportunities, while those with lower wealth may be unable to invest in such activities, leading to a widening of the wealth and income gaps. As those with little wealth struggle to increase their skillset they fall behind their better-off peers in terms of earnings and wealth accumulation, leaving them not only materially poorer, but also more vulnerable in the face of economic shocks, and less able to take advantage of opportunities. In time, this creates a vicious cycle that reduces the overall supply of talent in the economy hampering growth and prosperity for all.

Understanding the complex interactions between wealth, income and human capital is key to developing policies that are tailored to the needs of different groups and promote economic development and social mobility. For example, recent university graduates who are wealth-poor (due to student debt) but rich in human capital may require different policies than those who are more affluent at the moment but have a worse future outlook. Policies that focus on increasing wealth, such as tax breaks for the wealthy, may do little to address the unique needs of this group, as they already have high levels of human capital. Alternatively, policies that focus on increasing access to affordable education and training, or providing financial support to graduates struggling with student debt, may have a greater impact on this group's economic prospects.

Quantitative economic models that produce realistic distributions of human capital, income and wealth are well-suited to address the questions and challenges arising from the relationship between wealth, income and human capital. These models can provide insights into the relationship between these factors and help identify the best policy interventions to reduce wealth disparities and ensure that everyone has access to the resources and opportunities needed to build and maintain their human capital. Such models can also be used to quantify the effects of wealth disparities, economic shocks and other factors on human capital and wealth accumulation and the interactions and feedback loops between the two.

The heterogeneous agent incomplete market model based on the works of Aiyagari (1994), Bewley (1986), Huggett (1993) and Imrohoroglu (1989) has become the workhorse model of the economic analysis of wealth inequality and has been applied to study a vast range of economic questions (see e.g. Quadrini & Rios-Rull (2015) for a survey). A key feature of these models is that labour earnings are assumed to be exogenously given so that the households' problem focuses on the consumption-savings decision given a stochastic sequence of earnings. This means that most of these models do not consider the feedback effect that wealth inequality has on the distribution of earnings. Recent research has emphasized the role of family wealth for the educational outcomes of young adults. Scions from wealthy families are more likely to complete high school, enrol in university and complete their degree giving them a clear educational advantage in the labour market (see Lovenheim (2010), Karagiannaki (2017), Dräger (2022)). But, the link between wealth and human capital accumulation is not limited to (higher) education decisions. Both theoretical and empirical contributions to the study of human capital emphasize the importance of learning throughout the active working years (see Becker (1964), Ben-Porath (1967), Acemoglu (1998), Ma et al. (2020)). With longer working lives and increasing speed of technological change, such "on the job" or "post-formal education" training and skill accumulation is becoming increasingly relevant as evidenced by an increasing focus on "lifelong learning". Contrary to traditional higher education environments, financing for individuals interested in investing in such skills is much more limited making the role of pre-existing wealth more relevant for its study.

In this paper I provide an extension to the standard incomplete markets model, by allowing workers to invest in "risky" human capital in the spirit of Krebs (2003), resulting in an endogenously generated joint distribution of wealth, human capital and income. Specifically, I develop a general equilibrium model featuring endogenous wealth and human capital to illustrate some of the interactions between wealth and human capital. Labour earnings depend on human capital and idiosyncratic productivity shocks similar to Guvenen et al. (2014), giving rise to an endogenous earnings distribution. I calibrate the model to the UK economy in the pre-Covid-19 period using detailed household microdata from the UK's longitudinal household survey.

I analyse the interaction of wealth and human capital in the stationary equilibrium and find that there are important non-linearities in human capital investments. At low levels of human capital, when investments should yield high marginal returns, workers with low levels of wealth invest considerably less in accumulating human capital than their counterparts with more wealth. This pattern also holds true for workers with considerably higher levels of human capital. This suggests the existence of low-wealth poverty traps, where individuals with low wealth struggle to improve their skillset and therefore lag behind comparable individuals who have higher levels of wealth through previous savings or inheritance.

The steady state results provide evidence to the effect, that aside from ability and opportunity, wealth plays a key role in influencing the distribution of human capital and earnings. Further, it allows me to quantify these effects and assess the long-term impact of small initial differences in household wealth. I show these to result in quantitatively large differences in human capital accumulation that can persist for years and therefore entrench inequalities across generations.

I use the model to analyse the economic dynamics of the distribution of human capital in the aftermath of an unexpected economic shock and find a clear nonlinear pattern, with households close to the borrowing constraint exhibiting the strongest negative reaction to the shock. I show that this can lead to persistent increases in human capital and earnings inequality with consequences for future economic growth and development. This suggests that the initial distribution of wealth and human capital matters considerably for the aggregate response to an aggregate shock, with higher numbers of borrowing-constrained households leading to a bigger, more persistent drop in human capital and a slower recovery. Large recessions can be associated with persistent scarring effects, both for the individual and the larger economy (c.f. Ouyang (2009), Huckfeldt (2022)). The mechanism showcased by the model hints at a possible role for savings (or rather the absence of the same) in helping to explain these effects, but also in addressing other aggregate phenomena, such as the UK productivity puzzle (see Crafts & Mills (2020).

This highlights the need for public insurance policies to protect the most vulnerable in society during large economic shocks like the ones the UK has been experiencing in recent years. Providing support and insurance to low-wealth, low-income households is thus shown to not just be a charitable act to protect individual welfare, but rather an investment into the future stock of national human capital that has important implications for economic recovery and progress going forward.

I explore this point further by evaluating the medium-run impact of the Covid-19 pandemic through the lens of the model. The Covid-19 crisis in the UK combined a highly nonlinear economic shock, that affected workers across the income distribution in different ways (see Blundell et al. (2022), Marmot & Allen (2020), Stancheva (2021) and references therein), with a large-scale policy intervention by the government aimed at protecting lives and livelihoods across the UK. I calibrate a Covid-19 scenario, that captures both the highly unequal nature of the economic impact of the pandemic, as well as the emergency benefits intended to offset the loss of earnings and employment. This allows me to make model-based predictions of the medium-run impact of the pandemic along the lines of human capital, wealth and earnings inequality, as well as evaluate the effectiveness of the policy response and compare and contrast other possible intervention scenarios.

The analysis shows that Covid-19 is likely to lead to a persistent loss of human capital for the aggregate economy over the coming decades. While a strong intervention has prevented an even more substantial reduction, growth will likely slow as a result of a lower overall skill supply, exacerbating a decade of low productivity growth (Crafts & Mills (2020)).[1]

[1] While the analysis in this paper focuses on the skills of individuals of prime working age, there is a growing worry about the impact of the pandemic on the skill formation of children and young adults in full-time education (c.f. Blundell et al. (2022)). As these cohorts enter the labour market in the coming years and decades they will likely amplify these trends.

Comparing different policy interventions I show that in the absence of additional support measures low wealth households would have been most affected by the negative implications of the Covid-19 pandemic, leading to higher levels of human capital and wealth inequality. Successful policy intervention thus prevented welfare losses on the order of around 1.15% of lifetime consumption for the average household. An alternative UBI-type support system is shown to be slightly more welfare efficient but does not command a democratic majority amongst households.

The rest of this paper is organized as follows: Section 2 presents the model; Sections 3 and 4 present the calibration and the stationary economy; Section 5 showcases the response of the model economy to an unexpected productivity shock; Section 6 expands the analysis by analysing the medium run effects of the Covid-19 pandemic in the UK under different public insurance regimes; and Section 7 concludes.

2 Model

The economy is comprised of a continuum of infinitely lived household dynasties, distributed on the interval $I = [0, 1]$. Time is discrete and indexed by $t = 1, 2, 3, ...$ with each step corresponding to a year. At each point in time, the household is represented by a prime-aged worker (25-55 years) who supplies labour in a competitive labour market and makes decisions about consumption, savings and human capital accumulation during their working life. Households derive utility from consumption and receive income from labour earnings net taxes and transfers and the return on their savings which they supply to firms in a competitive capital market.

Human capital in the model refers to a unidimensional index that summarizes the total of all the worker's skills, talents and experiences that are relevant for general labour productivity. As the model focuses on prime-age workers, the model abstracts from higher and further education decisions. The accumulation of human capital in the model therefore refers primarily to skills that are acquired throughout the working life, through training or other means. Human capital accumulates based on the worker's human capital investment decisions and is subject to a small rate of depreciation over time.

A worker's labour productivity depends on two factors: their level of human capital and some exogenous productivity factor that captures how well their skillset is matched to the demands of the job, as well as other idiosyncrasies. Over time, the workers' productivity state does fluctuate stochastically creating uncertainty about future productivity. This means that human capital is *risky* in the sense that two workers with identical levels of human capital can have different levels of labour earnings depending on their exogenously assigned productivity state. Since the assignment of productivity states follows a random process, household dynasties differ in their history of labour productivity, leading to inequality in wealth, income and human capital.

Workers are subject to a risk of leaving the labour force ("dying"), which is realised through an exogenous random process. When workers leave the labour force, the household produces a new worker who steps into the fray.[2] The offspring inherits their parent's wealth[3] and part - but not all - of their human capital. This process creates intergenerational persistence in both wealth and human capital.

2.1 The households' preferences and constraints

The household dynasties desire to maximize their expected lifetime utility:[4]

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \tag{2.1}$$

where $\beta \in (0,1)$ is the discount factor and $c_t > 0$ is consumption in period t. The utility function takes the standard CRRA form:

$$u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}, \tag{2.2}$$

where $\sigma \neq 1$ is the coefficient of relative risk aversion. The CRRA form ensures that $u(c_t)$ is twice continuously differentiable, strictly increasing and strictly concave.

The household receives income from its asset holdings a_t in line with the risk-free interest rate r_t as well as labour income from selling its labour services in the labour market. These labour earnings are subject to a progressive income tax regime by the government which taxes earnings from high earners and provides some transfers to low-earning individuals. Net labour earnings are defined following Heathcote et al. (2017):

$$y_t = \tau_1 \left(l_t w_t\right)^{1-\tau_2}, \tag{2.3}$$

where τ_1 and τ_2 are two parameters describing the tax and benefits system; w_t is the

[2]To simplify the computation of the model I do not specify a full life cycle model, and instead stick with a "hybrid" model of stochastic "death" and replacement (see for example Rios-Rull (1996)).

[3]I also do not model bequests and inheritance dynamics explicitly, and instead assume that the bequest motive is sufficiently strong, that the terminal value function of a hypothetical finite horizon life cycle model is equivalent to the equivalent continuation value in the infinite horizon case. For a discussion of the implications of using a dynastic framework relative to a life cycle model see for example Hubmer et al. (2021). See Telemo (2022) for a life cycle model with bequests.

[4]I drop household-specific subscripts for ease of exposition.

wage rate and l_t is the worker's labour productivity, which will be described further in the next subsection. As described in Heathcote et al. (2017), $\tau_2 \in (0,1)$ governs the progressivity of the tax and transfer system[5], while $\tau_1 > 0$ sets the average level of tax. The government's sole objective is to redistribute income in this manner. Any shortfall in revenue is covered by the government in an unmodelled way; similarly, any excess tax revenue is spent in a way that does not affect firms or households.[6]

The households budget constraint is therefore as follows:

$$c_t + a_{t+1} + x_t = (1 + r_t)a_t + y_t, \tag{2.4}$$

where $a_{t+1} \in [a_{\min}, +\infty)$ denotes total end-of-period assets, which are subject to a borrowing limit[7] $a_{\min}$ and $x_t \geq 0$ denotes the worker's investment into his human capital.

2.2 Human capital accumulation

Workers accumulate human capital, by investing in a human capital investment good x_t the price of which is normalised to 1 for convenience. The investment good refers to expenditures that increase skills and do not count towards pleasurable consumption, such as training courses, professional certification exams or data science boot camps, but also activities that improve social and interpersonal skills, such as club memberships, and expenses to attend conferences, mentorship, and networking events.[8] More generally, x_t implicitly represents a utility cost that might be inherent in learning new skills, by absorbing resources that could have been used for consumption.

The household does not have an explicit budget of allocated time, but x_t can also be understood as a time cost, standing in for lost earnings that the household has to forgo in order to pursue skill improvement activities. In this case x_t stands for resources that are earned by the household, but then have to be paid back to their employer to cover training costs. [9]

Human capital accumulation is governed by the following law of motion, which is a

[5]If $\tau_2 = 0$ this corresponds to a flat tax rate of $1 - \tau_1$ while in the case of $\tau_2 = 1$ there would be complete redistribution.

[6]I assume this minimal government set up for simplicity. Having a more active fiscal policy, particularly in the context of training and education would be an interesting extension of the model.

[7]In principle, assets are unbounded above, but for the computational implementation of the model I use a grid with a finite upper bound on assets. For the calibrations, I consider the upper bound is chosen sufficiently large that agents never attain it.

[8]I abstract from the fact that some expenditures (e.g. books) might be both pleasurable (c_t) and skill-improving (x_t).

[9]Note that this is the reverse of the argument made in Heckman et al. (1998), but the argument works this way as well.

modified version of the standard human capital accumulation equation that is used in the literature (see Ben-Porath (1967), Heckman et al. (1998), Krebs (2003)):

$$h_{t+1} = \delta(\varkappa)h_t + \chi x_t^{\upsilon} \tag{2.5}$$

where $\delta(\varkappa) \in (0,1)$ is the depreciation rate of human capital and $\chi > 0$ and $\upsilon \in (0,1)$ are two parameters governing the transformation of the investment good x_t into human capital. The depreciation rate of human capital can take two distinct values, depending on whether the worker is exiting the labour force in that period. Whether the worker stays in the labour force or leaves is summarised by an indicator variable $\mathbf{1}_{\varkappa}$. If $\mathbf{1}_{\varkappa} = 1$, the worker ends his working life at the beginning of the period, and the household's human capital depreciates at rate $\delta(1) = \delta^{exit}$. A worker may decide to exit the labour force for a number of reasons, such as early or scheduled retirement, redundancy, or as a result of a severe illness or other impactful life event. The probability of such an event is given exogenously and has a probability ω. In this case, the depreciation rate effectively determines what share of the worker's human capital gets passed on to the next generation. If $\mathbf{1}_{\varkappa} = 0$ on the other hand, the worker remains in the labour force and the household's human capital depreciates at rate $\delta(0) = \delta^{stay}$.

I assume that workers cannot have negative levels of human capital (skill), so that $h_{t+1} \in [0, h_{\max}]$, where $h_{\max}$ is an upper bound determined by the limitations of human nature. In practice, however, the choices of functional form ensure that households will always choose values of h_{t+1} that lie within the boundaries of the interval, as long as $h_{\max}$ is chosen to be sufficiently large. Additionally, the assumption that workers "die" and lose a majority of their human capital helps to bound the problem here (see also Khun (2008) for a discussion of the case of permanent income shocks).

2.3 Labour productivity

As mentioned above, a worker's labour productivity depends on two factors: i) his stock of human capital h_t; and ii) factors exogenous to the worker that determine how productive his human capital is, e.g. his industry and occupation. These factors are summarised in the random variable e_t. Combining these, a worker's effective labour supply is given by:

$$l_t = e_t * \log(1 + h_t), \tag{2.6}$$

where $e_t > 0$ is the value of the exogenous productivity state, and h_t is the worker's current stock of human capital. The functional form $g(h) = \log(1 + h_t)$ is chosen to provide an increasing and concave function that transforms raw human capital into labour services.[10] This implies that exogenous labour productivity does not exclusively determine labour earnings, but rather is a contributing factor together with human capital (see for example Guvenen et al. (2014), for a similar mechanism). This also distinguishes the model from other models that include human capital and asset inequality (see Krebs (2003), Huggett et al. (2011)), where human capital is the sole ingredient generating labour income.

Exogenous productivity follows an $AR(1)$ process in logs:

$$\log(e_t) = \phi \log(e_{t-1}) + \varepsilon, \tag{2.7}$$

with $\varepsilon_t \sim N(0, \eta^2)$ and $|\phi| < 1$. Assuming this process for labour productivities is standard in much of the heterogenous agent incomplete market literature, and ensures that e_t follows a stationary process and therefore has an ergodic stationary distribution.

2.4 Aggregate quantities and market clearing

The economy has a single representative firm, which rents both capital and labour services from the workers in competitive labour and capital markets. The final output is produced by an aggregate Cobb-Douglas function:

$$Y_t = AK_t^\alpha L_t^{1-\alpha} \tag{2.8}$$

where A is total factor productivity and $\alpha \in (0,1)$ is the capital share of income. Capital depreciates at a constant rate $\zeta \in (0,1)$.

The representative firm's problem is given by:

$$\max_{K_t, L_t} AK_t^\alpha L_t^{1-\alpha} - w_t L_t - (r_t + \zeta) K_t \tag{2.9}$$

Assuming perfect competition, factor prices equal their respective marginal products:

$$w_t = (1-\alpha) A \left(\frac{K_t}{L_t} \right)^\alpha \tag{2.10}$$

$$r_t = \alpha A \left(\frac{L_t}{K_t} \right)^{1-\alpha} - \zeta. \tag{2.11}$$

In order to satisfy the necessary condition for the existence of an equilibrium with

[10] Most studies in labour economics consider logged earnings or earnings residuals as their object of interest making this a reasonable choice (see for example Mincer (1974)). Adding 1 is merely a rescaling of the human capital measure that ensures that labour earnings cannot be negative.

finite assets, I assume that the real interest rate lies in the interval $r_t \in (-1, \frac{1-\beta}{\beta})$ (see Acikgoz (2018) and references therein).

The government raises taxes and issues transfers to households according to the progressive tax and transfers schedule described in (2.3). The net tax take of the government is given by T_t. Apart from redistributing income, the government has a number of other spending commitments G_t (e.g. defence and infrastructure) that do not affect the utility of the household or the problem of the firm. I assume that the government can always find ways of spending extra tax revenue or can borrow costlessly in international markets without impacting the national economy. Hence the governments budget is perfectly balanced in every period:

$$T_t = G_t. \tag{2.12}$$

2.5 Stationary equilibrium

I assume that the economy is initially in a stationary equilibrium, in which changes to the asset and human capital positions of individual households cancel each other out, so that all relevant aggregate variables remain constant (see Hubmer et al. (2021)). In this equilibrium the aggregate capital stock and the effective labour supply are constant across all periods: $K_t = K_{ss}, L_t = L_{ss}$. This means that the factor prices and the agents' value and policy functions are time-invariant, as long as the economy remains in the stationary state. The standard definition of a recursive stationary equilibrium applies and is found when the quantities of labour and capital demanded by the aggregate firm are consistent with the quantities of capital and labour services supplied by the households, given the factor prices.[11]

Specifically, the stationary equilibrium of the model economy is defined as follows:

1. The representative firm solves its profit maximization problem, and factor prices equal their respective marginal products, $w_{ss} = (1-\alpha)A\left(\frac{K_{ss}}{L_{ss}}\right)^{\alpha}$, $r_{ss} = \alpha A\left(\frac{L_{ss}}{K_{ss}}\right)^{1-\alpha} - \zeta$.
2. Given w_{ss}, r_{ss} the agents solve the stationary version of the households problem, giving rise to a stationary distribution Ψ.
3. Capital and labour markets clear, $L_{ss} = \int l_i d\Psi(i), K_{ss} = \int a_i d\Psi(i)$.
4. The government's budget clears, $T_{ss} = \int (w_{ss}l_i - y_i) d\Psi(i) = G_{ss}$.

[11] For details see Appendix.

I solve the household's problem using dynamic programming. For this purpose, I rewrite the households' problem in its recursive form:

$$V\left(a_t, h_t, e_t\right) = \max_{\{c_t, a_{t+1}, x_t\}} \left\{ \frac{\left(c_t\right)^{1-\sigma}}{1-\sigma} + \beta E_t\left[V\left(a_{t+1}, h_{t+1}, e_{t+1}\right) \mid e_t\right] \right\} \tag{2.13}$$

$$s.t.$$

$$c_t + a_{t+1} + x_t = (1 + r_{ss})a_t + y_t$$

$$h_{t+1} = h_t\delta_t + \chi x_t^{\nu}$$

$$y_t = \tau_1(w_{ss}e_t \log(1 + h_t))^{1-\tau_2}$$

$$c_t \geq 0, h_t \geq 0,\ a_t \geq 0, x_t \geq 0$$

The main added difficulty of this problem is the existence of human capital as a second state in addition to assets. Models with two state variables have become increasingly common in macroeconomic modelling frameworks, but these tend to focus on a liquid and an illiquid asset class (see Kaplan & Violante (2014), Kaplan et al. (2018)). Human capital differs from physical assets in that agents can augment it, but cannot sell it off, implying an asymmetry in the possible margins of adjustment (see also Angelopoulos et al. (2021), who solve a model with health and assets).

I solve the household's problem using a modified version of the endogenous gridpoint algorithm for two state variables presented in Auclert et al. (2021). The main trick I exploit is the fact that the human capital accumulation function can be thought of as providing a convex adjustment cost function, which allows me to use some of the techniques developed in their paper.[12]

The main outcomes of interest are the value function $V\left(a_t, h_t, e_t\right)$ and the policy functions for assets, consumption and human capital investment, which I denote as $g^a\left(a_t, h_t, e_t\right)$, $g^c\left(a_t, h_t, e_t\right)$ & $g^x\left(a_t, h_t, e_t\right)$ respectively. The policy functions, together with the transition function for the exogenous productivity state e_t, and the retirement shock $\varkappa_t$ define a transition function Q^{ss} that maps current period states into next period states. The distribution Ψ is the stationary distribution associated with Q^{ss}.

I find the stationary distribution of all endogenous and exogenous variables, using a non-stochastic simulation (histogram) approach (see Heer & Maussner (2009), Young (2010), Angelopoulos et al. (2021)). For this I begin by creating the histogram of the joint state space across all endogenous and exogenous variables: $\tilde{\Psi}(a, h, e) \geq 0$ with total probability mass 1. I then define an approximation $\tilde{Q}^{ss}$ to the transition function Q^{ss} using linear interpolation to map next period states that fall between gridpoints back

[12]For details on the solution algorithm, see Appendix.

onto gridpoints (see Angelopoulos et al. (2021) for a description of the interpolation procedure with two endogenous state variables). Initializing $\tilde{\Psi}_0$ with uniform mass, I calculate the next period distribution as:

$$\tilde{\Psi}_{n+1} = \tilde{\Psi}_n \tilde{Q}^{ss}. \tag{2.14}$$

I continue updating $\tilde{\Psi}_n$ until the difference between successive histograms becomes less than a pre-set convergence criterion. I use the final histogram $\tilde{\Psi}$ from this procedure as an approximation of Ψ. Using enough grid points across the state space, this procedure can approximate Ψ to any desired degree of accuracy.

After describing the model and its solution method in this section, in the next section, I will describe the calibration procedure, and show that the model can capture key features of the empirical distributions of human capital, wealth and income in the UK.

3 Calibration

The calibration of the stationary equilibrium follows a mixed approach. I set a number of parameters to standard values that are used across a number of similar calibration exercises. Further parameters are set prior to solving the model using information specific to the economic situation that I study in this paper. All remaining parameters are set so that the model matches key features of the pre-Covid-19 economy in the UK over the period 2009 - 2019.

For this purpose I use a sample of prime-age working individuals from the UKLHS "Understanding Society" (UnSoc.) between 2009 and 2019. UnSoc is a comprehensive and representative panel survey and has been used to study a wide range of economic and social relationships. For this paper I focus on full-time, employed individuals aged between 25 and 55 years who are the heads of their respective households.[13]

3.1 Deep parameters

A number of parameters are set to standard values that are commonly used in the literature. These include the value of the coefficient of relative risk aversion σ in the household's utility function and a number of parameters related to the production function of the representative firm - α, A and ς.

[13] I further drop those individuals who are not observed in at least 3 consecutive periods and further trim the top and bottom 1% of gross labour earnings in each year. For more details on sample selection see Appendix.

I set the coefficient of relative risk aversion $\sigma = 1.5$ which is a standard value for the UK (see Angelopoulos et al. (2020) and references therein). I set total factor productivity $A = 1$ as a convenient normalisation and set the capital share $\alpha = 0.3$ (see Angelopoulos et al. (2020) and references therein) and the depreciation rate of capital $\zeta = 0.1$ to match the 10% annual depreciation rate used in Faccini et al. (2013).

3.2 Taxes and transfers

The UK has an active taxes and transfers policy which achieves a considerable level of redistribution - at least compared to the United States - from high earners to low earners (see for example Belfield et al. (2016), Blundell et al. (2018), Blundell & Etheridge (2010)). Evidence for this can also be found in my sample, as is presented in Figure 1 below. Figure 1 plots disposable income after taxes and transfers against market income for the UnSoc sample. As can be seen, a considerable portion of observations lie above the 45-degree line, indicating that these individuals receive more income through transfers than they pay in taxes. On the other hand, more affluent workers tend to make a net contribution to the welfare system. A progressive tax system has important implications for human capital accumulation. Progressive labour income taxes reduce the incentive to accumulate human capital, particularly for individuals who already have high earnings, as marginal tax rates decrease the return to high skills. On the other hand, transfers to the bottom of the earnings distribution provide valuable resources to low-income individuals who can use these resources to invest in skills. For these reasons, it is important to include a stylized tax and transfer system as described in equation 2.3 in the previous section.

To calibrate τ_1 and τ_2 I follow the strategy of Heathcote et al. (2017): Assuming that the tax and transfer system takes the form described in equation 2.3, I estimate the equation in logs using data on gross labour income and net labour income plus social benefits. The data estimates are $\hat{\tau}_1 = 4.01$ and $\hat{\tau}_2 = 0.41$ respectively. The second of these parameters can be used without modification, but since $\hat{\tau}_1$ governs the average level of taxes, I need to normalise it for the level of income in the stationary economy.[14] As noted in Heathcote et al. (2017) the break-even level of income using this modelling approach is given by $y^{breakeven} = \hat{\lambda}^{(\frac{1}{\hat{\tau}_2})}$, where $\hat{\lambda} = \exp(\hat{\tau}_1)$. Given the estimates, this break-even level of income is roughly £17,100 which is quite close to the annual salary of a worker working full-time on minimum wage. I calculate the ratio of this break-even income to the average income level in the sample resulting in a ratio of 0.577. I then set the calibrated value for $\hat{\tau}_1$ to match this ratio assuming an average income level of 1 in the stationary

[14] I set this to 1 as a convenient normalisation.

model economy, resulting in a calibrated value of 0.7976.

Figure 1: Representation of the UK's Tax/Transfer system for the UnSoc sample

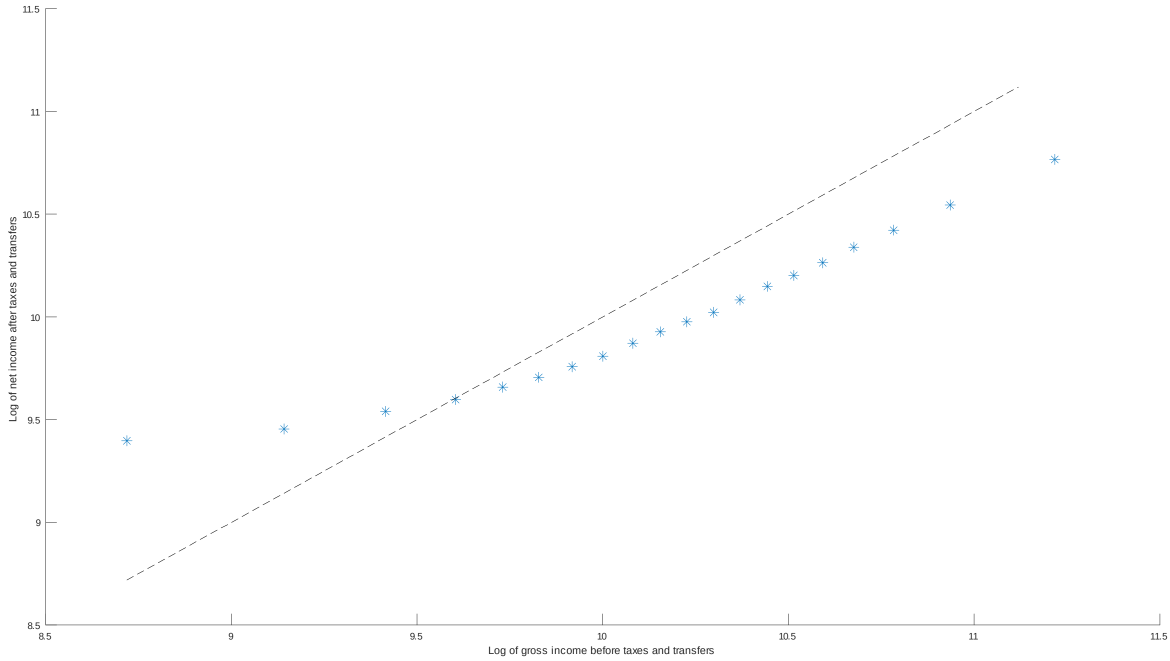

Note: Each datapoint represents the average in a quantile of the pre-redistribution income distribution. For details on variable construction and sample selection see this section and Appendix.
Source: Pooled UnSoc sample (2009 - 2019).

3.3 Stochastic labour force exit

The model features stochastic exit of the workers from the labour force. Every period workers exit with probability $0 < \omega < 1$ after which they are replaced by their household with a fresh worker. This approach is motivated by a desire to introduce turnover in the labour force into the model without adding the additional computational complexities of specifying a full lifecycle model. In my calibration, ω is set to $\frac{1}{30}$ to match the 30-year work life of the sample, and δ^{exit} is set as part of a minimum distance procedure to match some relevant data moments. Setting the expected work life to 30 allows me to focus on the savings and human capital investment behaviour of prime-aged workers. It is well understood that workers close to retirement have different incentives with regard to their savings and human capital investment behaviour (see Ben-Porath (1967), Huggett et al. (2012), Lachance (2012)).

An exit results in a worker being replaced by their offspring. An exiting worker can pass on all of their wealth but only part of their human capital to their child. This transfer of human capital is modelled by a higher level of human capital depreciation in the period of the death: $0 < \delta^{exit} < \delta^{stay} < 1$. The modelling choice in this case is driven by two considerations: 1. The model focuses on prime-age workers and hence abstracts from education choices made prior to labour market entry. It is likely that the children of parents with high levels of human capital have had more opportunities to accumulate skills prior to entering the labour market, for example by attending university. 2. There is a growing empirical literature that documents high intergenerational persistence of human capital across a variety of settings (see for example Black et al. (2005), Lundborg et al. (2018), Adermon et al. (2021)). Among other considerations, parents with different levels of income and/or education are likely to impart to their children different non-cognitive skills and attitudes, shaping their future labour market outcomes (see Acemoglu (2020)).

3.4 Data targets & model fit

The remaining parameters are calibrated as part of a minimum distance procedure to match a number of relevant data targets from the data. Targets represent some moments of the distribution of labour earnings, human capital and wealth in the data. For this purpose, I first purge all earnings and income variables of predictable variations by running a mincerian regression, controlling for a third-order polynomial of age, sex and region and year fixed effects. I retain the residuals from the Mincerian and use their exponent as my income measures. I also construct a proxy of human capital using measures of cognitive skills which are available in the data.[15] As UnSoc does not contain any information on wealth, I further use some standard values from the aggregate distribution of wealth.

I still require values for the exogenous productivity process e_t - namely the persistence ϕ and standard deviation of innovations η. As well as the parameters relating to the law of motion of human capital: δ^{stay} and δ^{exit} which are the values of the human capital depreciation for workers staying in or exiting the labour force; and χ and ν which are the linear and nonlinear parameters in the human capital production function, transforming the investment good into human capital.

I calibrate these remaining parameters, by minimizing the distance between a number of relevant data moments and their model counterparts. Specifically, I choose the mean and variance of the gross earnings distribution, the persistence of income, the capital-to-income ratio, the share of households with no assets and the variance and skewness of a proxy measure of human capital that is available in the data. For a detailed description

[15] For more details on variable construction see Appendix.

of all data targets, their model counterparts and their construction, see Appendix. I choose the parameters $\{\beta, \phi, \eta, \upsilon, \chi, \delta^{stay}, \delta^{exit}\}$ to minimize the sum of squared percentage deviations of the simulated moments from their targets. The whole set of target moments as well as the model fit for the calibrated model is presented in Table 1.

Table 1: Calibration Targets & Model Fit

Target	Data	Model	% deviation
Mean Gross Earnings	1	1.001	0.08 %
Variance Gross Earnings	0.288	0.277	-3.94 %
Persistence of Gross Earnings	0.887	0.888	0.14 %
% borrowing constrained	19	18.3	-3.70 %
Capital to Output Ratio	2.5	2.038	-22.64 %
Variance Human Capital	0.038	0.037	-3.39 %
Skewness Human Capital	-0.714	-0.715	0.08 %
Sum of squared percentage deviations	378.7324		

Note: UnSoc pooled sample and model simulation. For details see Appendix.

As can be seen, the final calibrated model matches the selected data moments very well. Particularly the mean, variance and persistence of gross income are well captured by the model, as are the second and third moments of the distribution of human capital. Figure 2 below highlights the fit of the calibrated model with respect to these two dimensions by plotting the cross-sectional distributions of gross labour earnings and human capital for both the UnSoc data and the model simulation.

Figure 2: Human capital and gross earnings in data and model

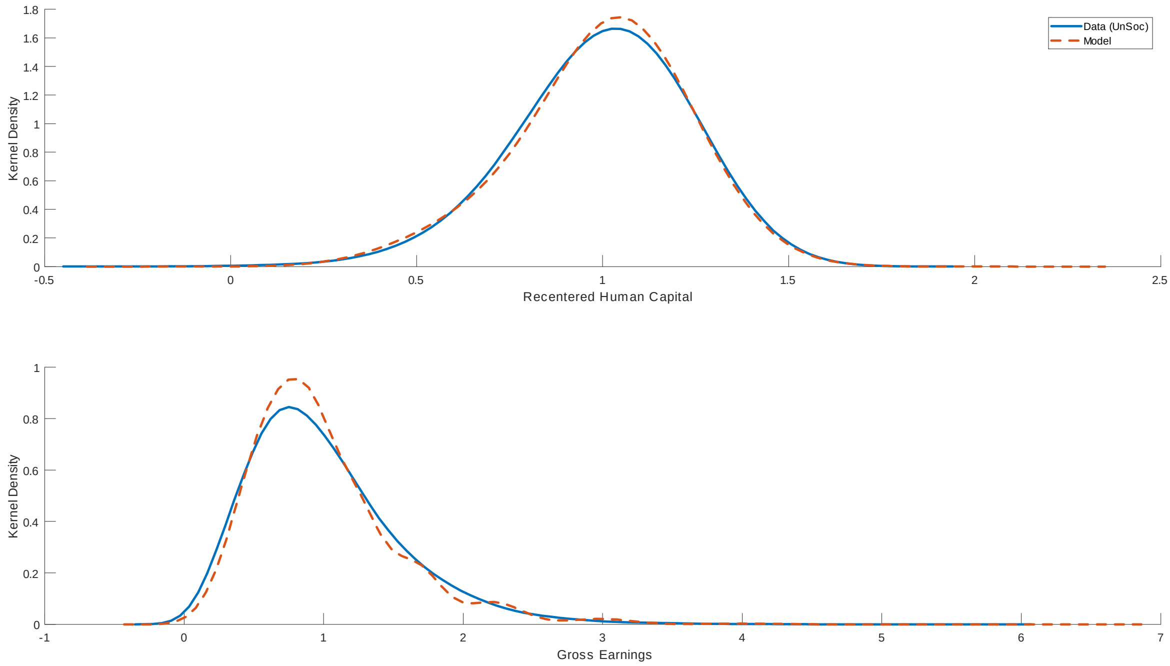

Note: Human capital measures are rescaled to mean 1 for comparability. Data gross earnings refer to exponentiated mincerian residual. For details on data construction see Appendix.
Source: Pooled UnSoc sample (2009 - 2019). and Model.

The main shortfall of the calibration is the low capital-to-income ratio which is about 20% below its target value. On the other hand, the left tail of the wealth distribution is matched quite well, with the model predicting a share of 18.3% borrowing-constrained households compared to 19% in the WAS data. For the purposes of this investigation, matching the overall capital stock is of secondary relevance, compared to capturing a realistic measure of wealth-poor workers.

Table 2 summarizes some key features of the stationary economy, providing an overview of key endogenous variables including moments that were not explicitly targeted by the calibration. As we have seen in the calibration section, the model captures the features of the UK earnings distribution quite well. Apart from earnings and human capital, the distributions of which are directly targeted by the calibration procedure, wealth is a key endogenous variable in the model. It is therefore important to evaluate the model fit with respect to household assets. The gini of assets for example is over 60 gini points, which is slightly below the empirical value of wealth inequality in the UK which lies around 0.71 (see Angelopoulos et al.(2020), (2021)) within the typical range for these types of models.

Table 3 provides a closer look at the model fit with respect to some important data moments. In addition to the UnSoc data which I used for the calibration, I further use information on the distribution of assets from the Wealth and Asset Survey (WAS), and consumption from the Living Cost and Food Survey (LCFS). To make the data comparable with the model output, I restrict the respective sample populations to households whose head is aged between 25 and 55 years.[16]

With respect to untargeted moments of the distributions of gross earnings and human capital the fit with respect to the data is quite good. The mean-to-median ratio, gini coefficient, CoV, and upper and lower tail measures of inequality are tracking the data quite closely, even though I only targeted the first and second moments or the second and third moments respectively. The fit with respect to the shares of earnings and human capital held by specific parts of the distribution is also quite well matched including for the shares of the top 10%.

With respect to assets and consumption, the model generally generates too little inequality. Particularly looking at measures of upper tail inequality, such as the P90-P50 ratio or the share of wealth and consumption of the top 80 or 90%, the model shows a lot less concentration of consumption or wealth. This might be a feature of the sample selection in these cases, but the difficulty of matching the concentration of wealth at the very top using HIM-type models is well-known.

Overall the model fit with respect to the targeted moments is quite good, particularly with respect to the distributions of human capital and earnings. In the next subsection, I will compare the performance of the model relative to a standard Aiyagari type and show that the human capital model does provide an improvement compared to standard methods.

3.5 Comparison with a standard Aiyagari model

To further put the stationary model results into context, I compare the model with human capital accumulation with an alternative version that does not feature endogenous skills. To allow for a fair comparison I set the exogenous earnings process of this standard Aiyagari model equal to the exogenous productivity process in the human capital model, and also standardise labour productivity so that mean gross income is equal to 1. All other parameters of the model are set to the same values as in the base model.

Table 4 presents a comparison of the distribution of earnings, wealth and consumption across both model versions.

[16] For more details on variable construction see Appendix C in Angelopoulos et al. (2021). I am grateful to Dr Spyridon Lazarakis for providing me with some of the summary statistics in this case.

In terms of mean earnings and assets both models produce remarkably similar results, but they do differ when we consider the distribution of these variables, as well as with respect to consumption in general. Earnings, both pre-and post-taxes and transfers are around 10% more unequal in the human capital model relative to the model with purely exogenous earnings. The inequality amplifying force of endogenous human capital extends to the case of assets which exhibit an around 10% higher gini coefficient. The difference is particularly noticeable along the lower tail, where the standard Aiyagari model showcases a much lower share of borrowing-constrained agents.

In the model presented in this paper, agents have to invest (at least in the long term) into their human capital, which creates additional demands on their resources. As a result, average consumption is about 15% lower in the human capital model relative to the model without endogenous skills. Additionally, consumption inequality is around 20% higher in the former model.

Table 2: Model Summary Statistics

	Gross Earnings	Net Earnings	Gross Income	Net Income	Human Capital	Assets	Consumption	HC Investment	Net Transfers
Mean	1.001	0.775	1.138	0.912	1.579	2.914	0.767	0.145	0.226
Median	0.914	0.756	1.023	0.865	1.617	1.686	0.732	0.141	0.157
Gini	0.274	0.164	0.275	0.193	0.105	0.609	0.172	0.212	0.65
Std.	0.527	0.232	0.592	0.32	0.304	3.616	0.243	0.056	0.3
CoV	0.526	0.299	0.52	0.351	0.193	1.241	0.317	0.39	1.325
$\frac{P10}{P50}$	0.516	0.677	0.5	0.634	0.697	0	0.672	0.52	-0.26
$\frac{P50}{P90}$	0.538	0.694	0.54	0.647	0.822	0.23	0.677	0.656	0.258
$Q_{<50}$	0.344	0.42	0.31	0.365	0.486	0.086	0.38	0.36	0.083
$Q_{>80}$	0.378	0.303	0.361	0.307	0.367	0.605	0.296	0.31	0.632
$Q_{>90}$	0.22	0.165	0.211	0.172	0.129	0.457	0.166	0.172	0.41

Note: Calibrated Model. Earnings refer to labour earnings only. Income refers to labour earnings plus returns on asset holdings.
Net Transfers are taxes paid minus transfers received. $Q_{<50}$ denotes the share of the bottom 50 percent.
Similarly, $Q_{>80}$ refers to the share of the population above the 80th percentile.

Table 3: Untargeted Moments

	Gross Earnings		Human Capital		Assets		Consumption	
	UnSoc	Model	UnSoc	Model	WAS	Model	LCFS	Model
$\frac{Mean}{Median}$	1.109	1.095	0.982	0.977	2.921	1.730	1.135	1.048
Gini	0.288	0.274	0.109	0.105	0.747	0.609	0.278	0.172
CoV	0.536	0.526	0.196	0.193	2.397	1.241	0.540	0.317
$\frac{P10}{P50}$	0.463	0.516	0.731	0.697	-0.060	0.000	0.524	0.672
$\frac{P90}{P50}$	1.894	1.859	1.209	1.217	7.456	4.348	1.910	1.477
$Q_{<50}$	0.299	0.344	0.424	0.486	0.016	0.086	0.308	0.380
$Q_{>80}$	0.367	0.378	0.249	0.367	0.742	0.605	0.368	0.296
$Q_{>90}$	0.214	0.220	0.129	0.129	0.548	0.457	0.219	0.166

Note: Comparison of calibrated model with model with selected data moments.
Source: Model; UnSoc pooled Sample (2009-2019);
Wealth and Asset Survey (WAS) (2014 - 2016);
Living Cost and Food Survey pooled sample (2014 - 2016).
For variable construction see Appendix and Angelopoulos et al. (2021).

Table 4: Comparison with standard Aiyagari model

	Gross Earnings		Net Earnings		Assets		Consumption	
	Aiyagari	HC model	Aiyagari	HC model	Aiyagari	HC model	Aiyagari	HC model
Mean	1	1.001	0.778	0.775	2.921	2.914	0.914	0.767
Median	0.901	0.914	0.75	0.756	1.969	1.686	0.893	0.732
Gini	0.25	0.274	0.149	0.164	0.547	0.609	0.139	0.172
Std.	0.477	0.527	0.212	0.232	3.194	3.616	0.232	0.243
CoV	0.477	0.526	0.273	0.299	1.093	1.241	0.254	0.317
$\frac{P10}{P50}$	0.59	0.516	0.733	0.677	0.04	0	0.718	0.672
$\frac{P50}{P90}$	0.59	0.538	0.733	0.694	0.269	0.23	0.738	0.677
$Q_{<50}$	0.429	0.344	0.505	0.42	0.139	0.086	0.403	0.38
$Q_{>80}$	0.571	0.378	0.495	0.303	0.606	0.605	0.272	0.296
$Q_{>90}$	0.344	0.22	0.277	0.165	0.376	0.457	0.151	0.166

Note: Comparison of calibrated model with model without human capital.
$Q_{<50}$ denotes share of the bottom 50 percent.
Similarly, $Q_{>80}$ refers to the share of the population above the 80th percentile.

The comparison in this subsection has placed the model in the context of the standard modelling frameworks of the HIM literature. The possibility to invest resources into skill accumulation amplifies inequalities in earnings, assets and consumption.

3.6 Discussion of calibrated parameters

The calibrated parameters are presented in Table 5 below. A main point of interest are the calibrated parameters related to the human capital accumulation function. As human capital is often difficult to measure in practice, empirical estimates of the parameters are rare and often context specific. It is however possible to contextualise the parameters with some comparable values used in other works, as well as simple common sense. Beginning with the depreciation rate during normal times δ^{stay}, the calibrated model suggests an annual depreciation of around 6%. In a recent study, Dinerstein et al. (2022) estimate a skill depreciation of 4.3% for Greek individuals who are out of work for the duration of a year. Assuming that even unemployed individuals invest a little bit into their skills, their estimate seems commensurate with a "raw" depreciation rate of 6%. Other evidence comes from more quantitative modelling exercises that are closer to this paper: Krebs (2003), whose work probably most closely mirrors this paper uses a depreciation rate of 6% which is in line with the calibration employed here. Huggett et al. (2011), estimate a lifecycle model using US data and find average human capital depreciation rates of around 2.5% using a "flat spot technique" for workers close to retirement, but it is important to note that since they use earnings as a proxy, their measure of human capital conflates exogenous productivity with human capital. Blundell et al. (2016), estimate a lifecycle model for women in the UK. They find depreciation rates between 5.7% and 8.1% annually, which fits neatly into the value found here. Finally, Fan et al. (2022) also estimate a lifecycle model on US data and find a baseline estimate of 8.6%, but argue the number should probably be lower.

Turning to the next parameter δ^{exit}, which measures the depreciation rate for workers who exit the labour force. Weighing in at roughly 25% this appears rather large, but given that it signifies the beginning of a new generation of workers it does not seem particularly excessive. There is little consensus on the exact magnitude of intergenerational transmission of human capital, other than broad agreement that it can be sizeable (see for example Black et al. (2005), Lundborg et al. (2018), Adermon et al. (2021)). Clark (2014) finds intergenerational correlation coefficients of between 0.7 and 0.8, which would be in line with the 0.75 suggested by the calibration, however, these numbers are disputed and other estimates suggest that values of around 0.45 are more appropriate (see the discussion in Solon (2018)). This suggests that the calibrated value might be slightly too low, but given the large degree of uncertainty and the very specific nature of the application, it is an acceptable value.

Finally, the nonlinear term in the function transforming the investment good into human capital. The calibrated value of 0.825 is comfortably below 1 and fits into the interval $(0.5, 0.9)$ which is suggested by Huggett et al. (2011) as reasonable values. In the same paper, they suggest that values around 0.65 provide the best model fit for their application, but it is worth pointing out that their law of motion is slightly different from mine, as well as the differences in institutional setting and sample. Krebs (2003) uses a value of 1, putting my calibration somewhere in between their values. Guvenen et al (2014), also calibrate ν to 0.8 in their baseline application.

Table 5: Calibrated Parameters

Parameter	Function	Calibrated Value	Method
ϕ	Persistence of Prod. Shock	0.882	Minimum Distance
η	Std of Prod. Shock	0.102	Minimum Distance
τ_1	Linear Taxes	0.798	Outside Calibration
τ_2	Progressivity of Taxes	0.411	Outside Calibration
δ^{stay}	Depreciation of HC	0.939	Minimum Distance
δ^{exit}	Intergenerational Transmission of HC	0.746	Minimum Distance
χ	Linear HC Production	0.531	Minimum Distance
υ	Decreasing returns in HC	0.825	Minimum Distance
β	Discount Factor	0.95	Minimum Distance
α	Labour Share	0.3	Standard Value
σ	CRRA Coefficient	1.5	Standard Value
ζ	Depreciation rate of capital	0.1	Standard Value
a_{min}	Borrowing limit	0	Natural Borrowing Limit
ω	Probability of labour force exit	0.03333	Outside Calibration
A	Total Factor Productivity	1	Normalisation

4 Stationary Economy

Having solved the model and calibrated it to match key features of the UK economy between 2009 and 2019, this section analyses the stationary distribution in more detail. The main focus of the analysis is on exploring how wealth and human capital interact and shape the incentives and constraints of the households. I analyse how wealth (or lack thereof) enables (prevents) households to build up human capital quickly and thereby (not) taking advantage of profitable labour market opportunities and how this mechanism generates large persistent differences in human capital and ultimately income and consumption between households. This analysis is then extended in the following section, where I analyse the behaviour of the economy in response to an unexpected aggregate productivity shock, which lays the groundwork for the application to the Covid-19 crisis in the penultimate section.

In the rest of this section, I explore some properties of the stationary economy, focussing on the agent's optimal decision rules and how these are both influenced by and contributing to wealth and human capital inequality.

4.1 Policy function analysis

In the model workers invest in both physical assets and human capital, implying the existence of trade-offs. Wealth provides a constant income stream due to interest payments and can be used to smooth consumption in the case of a drop in income. Human capital on the other hand provides generally larger returns than savings, but the return is risky since it depends on the stochastic productivity state e_t.[17] Furthermore, human capital is inalienable and cannot be used directly to smooth consumption in response to income shocks.

Figure 3 plots the policy functions of agents with different levels of wealth and human capital. The first row shows the level of investment into an agent's human capital $g^x(a_t, h_t, e_t)$, by their current human capital and wealth, as well as by different exogenous productivity levels. Moving from left to right, the subplots represent the policy functions conditional on different values of the exogenous productivity state e_t - namely the minimum, mean and maximum values of e_t. As the marginal product of human capital depends on the exogenous productivity state, we expect workers with higher human capital productivity to invest more resources into acquiring additional human capital. Indeed this is confirmed by comparing the levels of the different conditional policy functions. At the same current level of human capital and asset holdings, workers in higher productivity states, invest more into their human capital, in line with predictions based on their increased marginal return from doing so.

Slightly more unexpected is a striking pattern of nonlinearity emerging from these policy functions, which exist *conditional* on a given level of exogenous productivity. Nonlinearities along this dimension indicate differing incentives and constraints based on different values of human capital and wealth. With an increasing and concave function, mapping human capital into earnings, investment into human capital should be highest for low levels of human capital and then decrease in line with human capital's marginal product.[18] However, since workers with low levels of human capital also have low levels of labour earnings, their available resources are potentially quite restricted. Here wealth enters as a crucial factor, relaxing the budget constraint, so that human capital investment can be maintained at the desired level. This effect becomes obvious, particularly if one focuses on low levels of human capital. As a general trend, human capital investment x_t decreases as the initial level of human capital increases - a pattern that is consistent with the decreasing marginal returns of human capital due the concavity of $\log(1 + h_t)$. However, at very low levels of h_t workers with low levels of wealth invest much less into their human capital, than their more prosperous counterparts.[19]

[17] Since the return to human capital is risky, its average return must be greater than that of the riskless assets, or no household would ever invest in human capital. See also the general result and a discussion of this point in Krebs (2003).

[18] It might be useful to imagine the case of a perfectly frictionless version of the model as an example. In that case, investment in human capital should follow a "waterslide" type shape.

[19] In empirical studies the role of credit constraints on human capital accumulation can often be sub-

This simple graphical analysis shows how the borrowing friction in the model creates an interdependency between wealth and human capital, thereby making wealth inequality relevant for human capital accumulation and thus future earnings inequality.[20] It further shows that this interaction is much more impactful at low levels of human capital. At high levels of human capital, and across different productivity states, the policy functions are much more similar across the asset dimension than is the case at lower levels of human capital. This is likely driven by two factors: i) higher human capital implies higher labour earnings, and therefore even wealth-poor households are likely able to finance their investment plans out of current earnings ("income effect"); ii) as the marginal product of human capital decreases, investment plans drop so that they can be easier attained by resource-constrained workers ("price effect").

These observations are reinforced by looking at the changes in the worker's net asset positions $\triangle a_{t+1} = g^a(a_t, h_t, e_t) - a_t$ in the second row of Figure 3. Workers with low levels of human capital, but high levels of wealth draw down their assets to finance not only consumption but also a rapid expansion of their human capital. Agents close to the borrowing constraint on the other hand see very little change in their asset position until they have built up a reasonable amount of human capital.

stantial (see Lochner & Monge-Naranjo (2012) for a survey).

[20]This point will be explored further below in this section.

Figure 3: Policy functions for endogenous variables

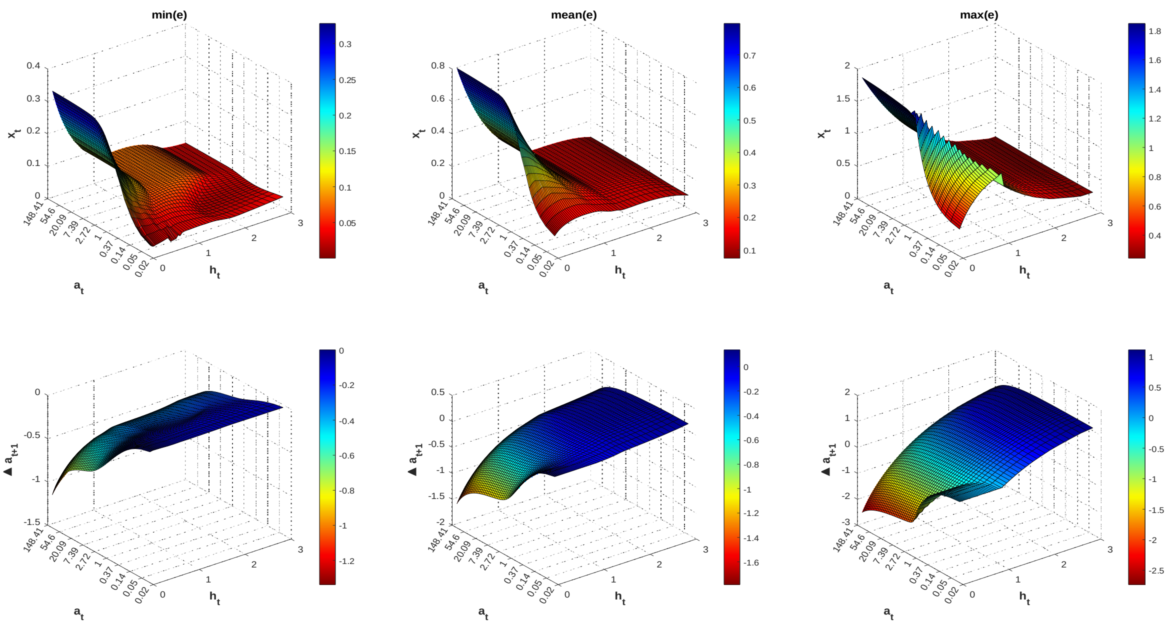

Note: Policy functions of endogenous variables for minimum, mean and maximum level of productivity. Assets are plotted on log scale.

4.2 The role of initial conditions

As the analysis of the policy functions has shown, wealth enables workers to accumulate human capital more quickly which has important implications for future earnings. Forward-looking, rational workers in this model make taking into account future expected costs and benefits of their present actions. Actions taken by the worker today affect their position in the next period, leading to a dynamically changing set of incentives and constraints. The value and policy functions that constitute the solution of the model encode this information in a way that is accessible to the agent, but not necessarily the economist trying to glean an insight.

To give a more detailed account of the interaction between wealth and human capital inequality in this model, this subsection provides some further analysis of how these interactions play out dynamically across time. In order to assess these dynamic effects, I simulate hypothetical households which differ only with respect to their initial level of wealth and track relevant measures of their human capital and earnings potential over time.

Figure 4 plots the evolution of human capital and gross earnings for workers starting with assets equal to the 20th, 40th, 60th and 80th percentile of the stationary distribution, conditional on different initial levels of human capital and the *same* level of initial productivity. I present the results normalized relative to a comparable worker with assets equal to the median of the stationary distribution of assets. This normalisation is provided for convenience, as in the long run the expectation of all households will converge to the unconditional average irrespective of initial conditions. This *ergodic* behaviour of the system means that across all households there is a small "drift" towards the population average, however, in the short and medium run, interesting trends in divergences emerge.

Specifically, each line in the top row of Figure 4 corresponds to the sequence:

$$\frac{E_0(\{h_t\}_0^T | h_0^{pc}, a_0^{pc}, e_0) - E_0(\{h_t\}_0^T | h_0^{pc}, a_0^{50}, e_0)}{E_0(\{h_t\}_0^T | h_0^{pc}, a_0^{50}, e_0)}, \tag{2.15}$$

where h_0^{pc}, a_0^{pc} indicate initial conditions equal to specific percentiles of the stationary distributions of human capital and wealth respectively. So for example, the lines in the first subplot of Figure 4 are obtained by setting h_0^{pc} to the 20th percentile of the human capital distribution, and then calculating the expected mean percentage deviations, setting in turn a_0^{pc} to the 20th, 40th, 60th and 80th percentile of the stationary distribution of assets. The subplots in the second row are obtained similarly, using gross earnings instead of human capital.

Turning to the subplots in row one, the most striking outcome is a large negative gap that opens up between workers starting with comparatively low assets and comparable workers with more savings. A worker starting close to the borrowing constraint at the 20th percentile of wealth, and human capital, has after 10 years accumulated about 10% less human capital than his twin who started with median wealth, even though both started with the same amount of human capital in the first period. This 10% gap in human capital translates into a sizeable effect on labour earnings, with a mean difference in gross earnings of roughly 6%. Compared to this a worker at the 40th percentile of assets has only lost just over 1% in human capital relative to the median wealth household. This nonlinear effect is emphasized by looking at the comparatively small differences between workers starting with the 60th and 80th percentiles of wealth, compared to workers at the 60th and the 40th or the 40th and the 20th percentiles.

This pattern of very strong negative effects for low-wealth households is consistent across different initial levels of human capital, even though it is attenuated somewhat at higher levels of h_t. For example, the difference between the 20th and 50th percentile of wealth households reduces to around 8% and 6% if we consider workers that start with human capital equal to the 50th or 80th percentiles of the stationary distribution respectively.

These observations confirm some of the earlier analysis: i) the resource constraint binds particularly strongly for households close to the borrowing constraint; and ii) at higher levels of human capital, the impact of initial differences is weakened somewhat.

An additional, striking feature, which this graphical analysis has uncovered, is the long persistence of these differences in initial conditions. Comparing for example 40th and 60th percentile households relative to the median wealth household, we see that the differences in terms of human capital and gross earnings are initially small, but over time small differences accumulate and a gap opens up. This hints at the existence of very persistent impacts of small differences in initial wealth as well as a strong role in the intergenerational transmission of human capital mediated by asset wealth.

Figure 4: Mean differences in human capital accumulation and gross earnings, by wealth quantiles

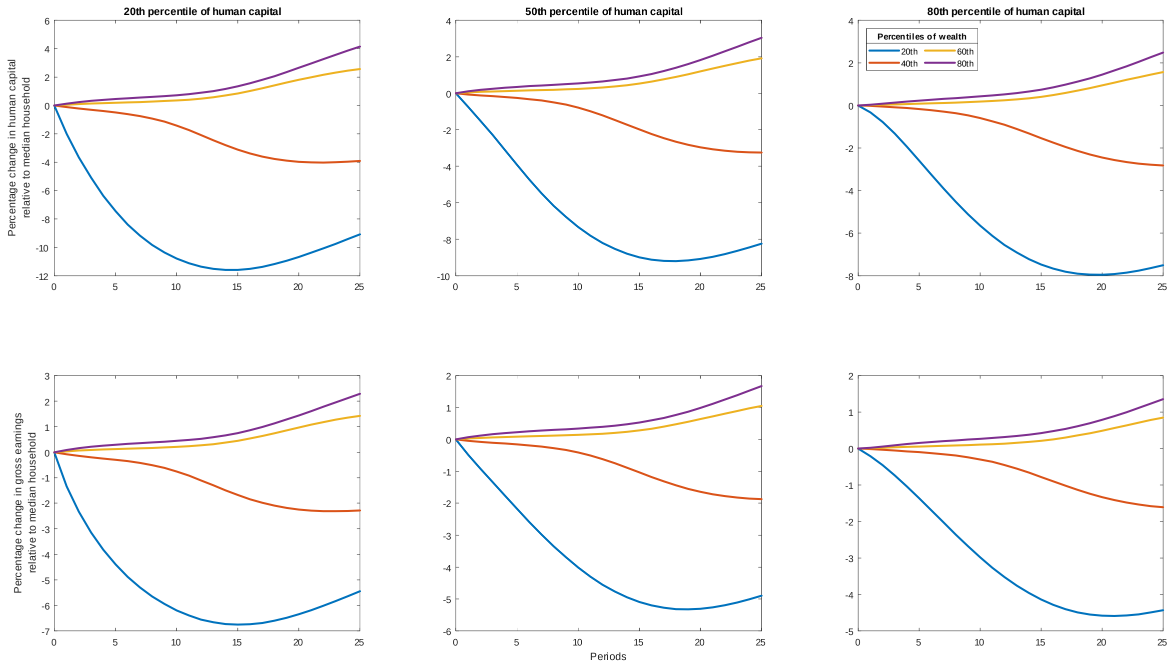

Note: Workers begin with mean level of productivity in the initial period.
All series are normalized relative to a household with 50th pctile of wealth in the initial period.

Focussing on average outcomes can sometimes understate another important function that wealth fulfils in the context of the incomplete markets model, namely helping the household weather bad shocks, and reduce the negative consequences of risk by providing a form of self-insurance. The workers in the model face risk in the form of exogenous productivity shocks. The standard analysis in the context of the HIM model suggests that savings allow households to smooth consumption across periods with high and low productivities. The model presented in this paper adds an additional dimension to this, as workers not only strive to smooth consumption but also maintain human capital at an optimal level. This adds to the agents' precautionary savings motives and provides an additional avenue for initial differences in wealth to have long-term consequences.

Figure 5 provides an illustration of the impact of initial differences in wealth on future human capital and by extension earnings risk, by plotting the variance of logarithms of human capital and gross earnings for different initial levels of wealth and human capital in a similar manner to Figure 4 above, only that I am plotting absolute values instead of percentage differences. For a household starting at $t = 0$ the value of the coefficient of variation with respect to some outcome at some time t in the future provides a summary measure of the distribution of the different realizations of that outcome that the household might face - hence it provides a measure of risk with respect to that outcome.

Focussing on the first row in Figure 5, with respect to human capital, higher levels of risk indicate more uncertainty about future earnings and available resources. Naturally, as all households begin from the same initial level of human capital, the variance of logarithms begins at 0 and then slowly rises as the exogenous shocks accumulate over time. Accordingly, an increase in human capital risk is expected, but it is notable that the rate of increase is largest for the households with the lowest levels of initial wealth. For example, for all households that began at the median level of human capital, after 10 years the expected distribution of human capital is about 300% more unequal for those with initial wealth equal to the 20th percentile, compared with their counterparts that started with assets equal to the 40th, 60th or 80th percentile. Again, the effect is highly nonlinear, with the difference between the 20th and 40th percentile households, far exceeding the difference between the 60th and 80th percentile households. This suggests that the poorest households are most exposed to human capital and therefore future income risk, as a result of their inability to use wealth as a tool to maintain human capital levels through times of low productivity. As before, higher initial levels of human capital help reduce the importance of these initial differences in wealth somewhat, but differences remain quantitatively large.

For a risk-averse household, the variance of future earnings is an important factor in determining welfare and well-being. The second row in Figure 5 provides an analysis of the future earnings risk that households face, based on their initial level of wealth. Overall the patterns are more similar compared to the case of human capital, as earnings risk is partly a function of the exogenous productivity process which is the same for all households. But still, after a couple of periods, there is a notable elevation of earnings risk for households with lower initial wealth which is a consequence of higher human capital risk.

In a standard incomplete markets model, the main function of precautionary savings is to enable consumption smoothing and thus counteract the negative consequences of income risk. By including human capital in the model, agents gain the ability to directly affect the stream of future incomes. This gives wealth an important additional function effectively allowing rich households to directly affect their future income risk. This potentially increases the value of savings for households and emphasizes the importance of accounting for wealth inequality when assessing income and earnings inequality.

Figure 5: Dynamic risk in human capital and earnings, by wealth quantiles

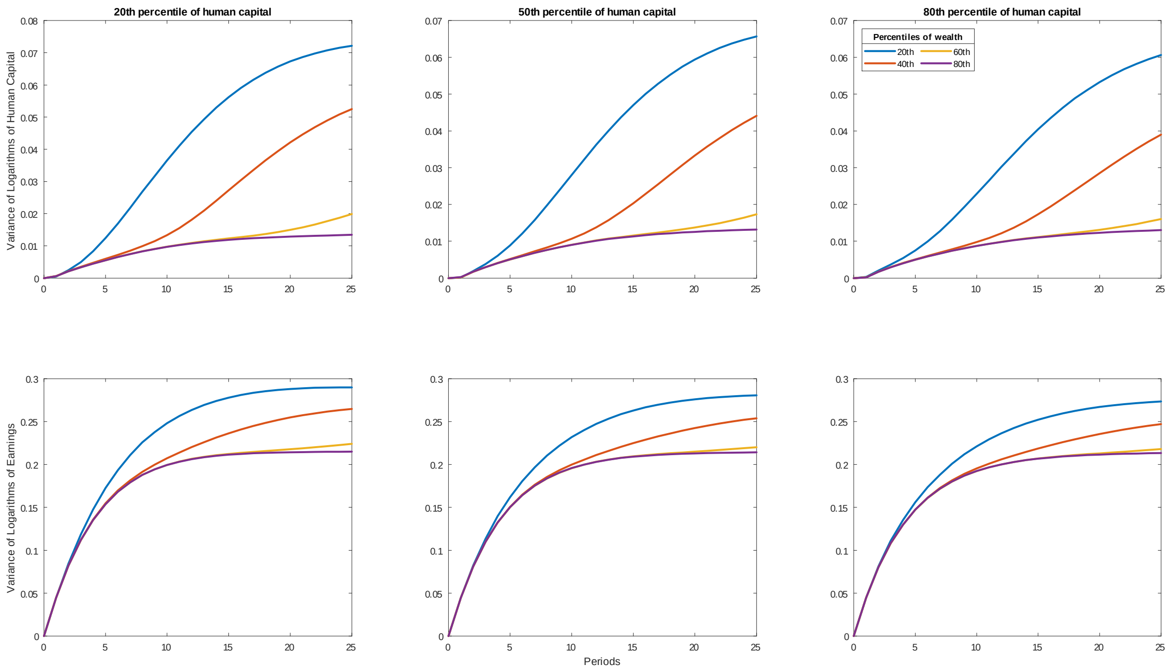

Note: Workers begin with mean level of productivity in the initial period.

The preceding analysis has shown, that initial differences in wealth can lead to large and persistent differences in human capital accumulation by otherwise identical households. But the relationship between wealth and human capital is not limited to flow in this direction. The perhaps more intuitive relationship between the two quantities suggests that workers with more human capital obtain higher earnings and are thus able to accumulate more wealth.

Figure 6 quantifies the effect of initial differences in human capital on wealth accumulation. The logic of the sequences plotted in Figure 6 is similar to those plotted in

Figure 4, with the difference that the dimensions of wealth and human capital are flipped - we are following households that start with the same amount of initial assets but with different endowments of human capital.

The plots reveal a clear wealth inequality enhancing role for human capital. Across different initial levels of wealth, households with lower initial human capital accumulate wealth much slower than their more skilled counterparts. These differences can be substantial: for households close to the borrowing constraint, those with human capital equal to the 20th percentile will after 10 years only have accumulated 60% of the wealth of a household that started with the median amount. In contrast, highly skilled households starting with the 80th percentile of human capital will have accumulated 20% more assets than the median household. These effects are large and highly persistent, even though the absolute difference does diminish as the initial level of assets increases.

Taken together with the previous analysis suggests that the model generates an important feedback loop from wealth to human capital back to wealth. Wealth and human capital inequality are therefore mutually reinforcing, suggesting that any analysis that aims to understand one of the two should not disregard the other.

Figure 6: Mean differences in wealth accumulation, by human capital quantiles

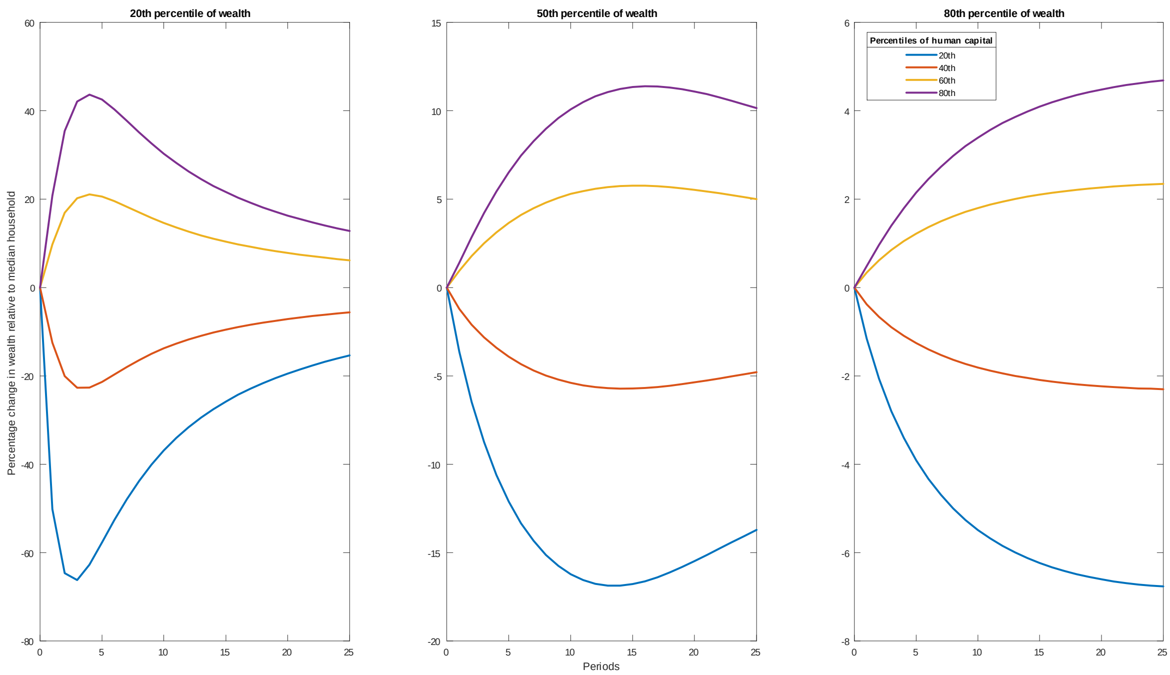

Note: Workers begin with mean level of productivity in the initial period.

In this section, I presented an analysis of the policy functions of the model in the stationary equilibrium. I showed that workers with higher human capital productivity optimally invest more resources into acquiring additional human capital. Furthermore, non-linearities in the policy functions emerged, indicating differing incentives and constraints based on different levels of human capital and wealth.

Subsequently, I presented an analysis of the importance of wealth, by analysing the dynamic behaviour of hypothetical households with different initial conditions. This analysis showed that workers starting with low assets accumulate up to 10% less human capital than their twins who started with median wealth, which translates into a sizeable effect on labour earnings over time. Moreover, these differences were found to be highly persistent, leading to potentially large differences in intergenerational inequality in human capital and income.

Further, I discussed the role of risk in the model. Higher levels of risk indicate more uncertainty about future earnings and available resources. I find, that the rate of increase is largest for the households with the lowest levels of initial wealth. Meaning that the poorest households are most exposed to human capital and therefore future income risk and as a result face the most uncertainty about their future consumption.

Finally, I illustrated the effect of human capital on wealth, showing that amongst households with the same initial level of wealth, higher-skilled households accumulate assets much faster leading to a positive feedback loop between wealth and human capital.

The analysis in this section has provided an insight into the dynamics of human capital accumulation in the model, and how initial differences in wealth can lead to persistent differences in human capital and consumption across households and vice versa. In the next section, I will complement this analysis by looking at the reaction of the economy to an unexpected aggregate shock. This approach allows a more detailed analysis of how different initial conditions interact with aggregate shocks to shape the dynamics of the economy and to show how differences in wealth interact with human capital inequality.

5 Aggregate dynamics

The last section analysed the policy functions of the model in the stationary equilibrium, showing that differences in human capital and wealth can have an impact on incentives and constraints, leading to an interaction between wealth and human capital inequality. I also looked at the importance of initial conditions, showing that even small differences in wealth can lead to persistent differences in human capital and earnings across households, further finding that the poorest households are most exposed to human capital and future income risk. In all this it was assumed that the economy is in a steady, stationary state and not subject to aggregate disturbances, large or small that affect the economic fortunes of all households at the same time.

However, in reality, the economy is rarely in a state of perfect stationarity. At the aggregate level, unexpected economic shocks are an unavoidable reality. These shocks can arise from a variety of sources and can have a significant impact on the economy. For the UK, some notable recent examples of economic shocks include the Great Financial Crisis of 2007/08, the Brexit referendum, the Covid-19 pandemic, and the war in Ukraine with its associated energy and cost of living crisis.

In this section, I will analyse the response of the model economy to an unexpected, one-off, aggregate productivity shock.[21] This analysis complements the steady state analysis in the previous section and also lays the groundwork for the evaluation of the Covid-19 pandemic in the following section.

5.1 Dynamic equilibrium

To implement the shock, I assume that the economy is in its stationary equilibrium, when at time $t = 0$, an unexpected productivity shock hits the economy, reducing total factor productivity by 1%.[22] The shock lasts for one period and then disappears, never to return again. The households observe the shock and perfectly forecast the path of all aggregate variables in the future and make their decisions based on their knowledge on the future evolution of factor prices.

This defines a perfect-foresight transition path for the economy.[23] Specifically, such a path for the economy is characterised by:

1. A time horizon $T > 0$ large enough that if the economy is hit by an aggregate shock at $t = 0$ it will have settled back into its stationary state by time $t = T$.
2. A path for the aggregate variables $\{K_t, L_t, A_t, G_t\}_{t=0}^{T}$.
3. Given $\{K_t, L_t, A_t\}_{t=0}^{T}$, the representative firm solves its optimization problem, and factor prices are given by $\{r_t, w_t\}_{t=0}^{T}$.
4. Given $\{r_t, w_t\}_{t=0}^{T}$, the households solve the dynamic version of their decision problem, and make choices consistent with optimal behaviour given the path of the aggregates.

[21] In general these are called "MIT" shocks by the literature.

[22] Interestingly, unlike in the case of a representative agent model, the response of the economy is not symmetric across positive and negative shocks. I do not analyse the response to a positive productivity shock in detail, as I want to focus on negative shocks in line with my empirical application, but the Appendix provides some illustrative figures.

[23] For details of the implementation, see Appendix.

5. Markets clear in every period.

6. The government's budget is balanced in every period.

The households optimization problem is now defined by a finite horizon problem, where the terminal value function at time T is equivalent to the value function of the stationary problem. The Bellman equation of this problem is defined as follows:

$$V_t\left(a_t, h_t, e_t\right)=\max_{\left\{c_t, a_{t+1}, x_t\right\}_{t=0}^T}\left\{\frac{\left(c_t\right)^{1-\sigma}}{1-\sigma}+\beta E_t\left[V_{t+1}\left(a_{t+1}, h_{t+1}, e_{t+1}\right) \mid e_t\right]\right\} \tag{2.16}$$

$$\begin{gathered}
s.t. \\
c_t+a_{t+1}+x_t=(1+r_t)a_t+y_t \\
h_{t+1}=h_t\delta_t+\chi x_t^\nu \\
y_t=\tau_1(w_t e_t \log(1+h_t))^{1-\tau_2} \\
c_t\geq 0, h_t\geq 0,\ a_t\geq 0, x_t\geq 0 \\
V_T\left(a_t, h_t, e_t\right)=V_{ss}\left(a_t, h_t, e_t\right)
\end{gathered}$$

I solve the household's problem by backward induction using a modified version of the algorithm used to solve the stationary household problem.

5.2 Aggregate effects

Figure 7 plots both the average and the distributional effects of the MIT shock on human capital investment, human capital, gross earnings, consumption and wealth. Investment into human capital drops sharply in the period of the shock, falling by about 0.4% relative to its steady state value. This is driven by both, a tightening of the households' resource constraint ("income effect"), as well as a reduction of the effective return to human capital going forward ("price effect").[24] Human capital investment recovers quickly thereafter, reaching a value of 99.9% of its steady state level 5 years after the shock. The drop in investment, affects the aggregate stock of human capital, although the reaction is much more sluggish. Human capital reaches a low of 0.05% below steady state after around 7 years, after which a protracted recovery begins. Human capital inequality follows suit, peaking at around 0.05% above its initial level. Pre-tax labour earnings drop by 1% on arrival of the shock and then mostly recover, but then remain around 0.1% below the initial

[24]Note that this is entirely driven by the endogenously generated wage rate, as the shock itself has no persistence.

level for a prolonged period. Similarly earnings inequality spikes following the shock, and then remains elevated for around 40 years. This suggests that human capital in the model can generate an internal propagation mechanism, that leads to a persistent "scarring" effect of recessions in terms of both average earnings and inequality. Consumption falls on average, but only slightly by 0.15%, as households draw down asset savings to protect themselves. Assets themselves fall by about 0.2%, and then recover slowly together with average consumption levels. Consumption inequality exhibits a considerable spike in the period of the shock, likely as a result of borrowing-constrained households being unable to borrow to bridge their temporary earnings shortfall.

Figure 7: Response of aggregate variables to an unexpected 1% drop in TFP

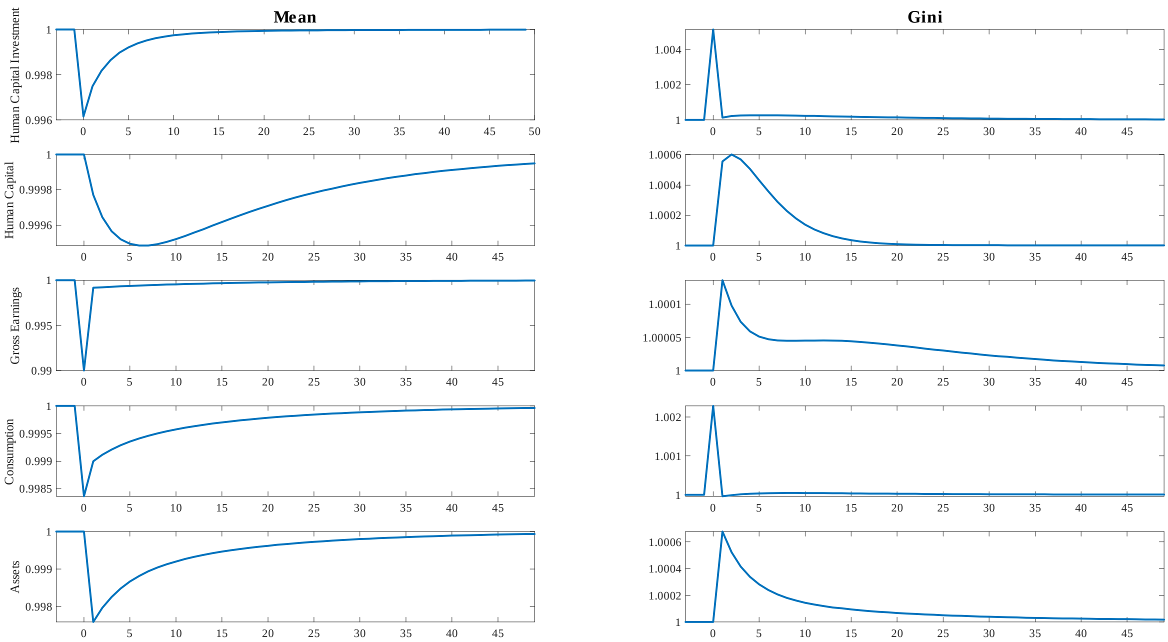

Note: All values are normalised to 1 in the stationary equilibrium.

The deviations analysed above must necessarily appear small, given the size of the MIT shock. However, what is remarkable is how even small shocks can have long, persistent impacts on human capital and earnings. To provide a better evaluation of the impact that these shocks can have over the lifetime of the agents, Table 6 below provides elasticities of a number of key endogenous variables of the model over different time horizons. Elasticities are calculated as the sum of percentage deviations from steady state divided by the magnitude of the shock. So for example, for endogenous outcome z the elasticity with respect to time horizon T is calculated as follows:

$$E_z^T = \frac{\sum_{t=0}^{T} \frac{z_t - z_{ss}}{z_{ss}}}{\frac{A_0 - A_{ss}}{A_{ss}}}, \tag{2.17}$$

where the subscript ss denotes steady state values.

Focussing first on the elasticities of the means, it can be noted that the impact effects of the shock are initially small, except for the case of earnings, which are directly affected by the fall in total factor productivity. Human capital and wealth stocks are fixed at the beginning of the period, so the shock has no immediate pass-through to these quantities, but even human capital investment reacts fairly inelastically, with an elasticity of around 0.39, whilst consumption has an even lower impact elasticity of 0.16. Looking at the medium-run impact of the shock, the 5 and 10-year elasticities the impact becomes more pronounced. Taken over a 5-year horizon, human capital investment exhibits a slightly more than unit elastic response, while the elasticity of gross earnings goes up to 1.3 suggesting that around an additional 30% of the initial shock has been added to earnings losses over that time period. Human capital and assets only respond with a lag to the shock, as was also indicated by Figure 6 above. After 10 years the elasticity of human capital has grown to around 0.4, while that of assets exceeds 1.3. In the long run, elasticities continue to grow, albeit at a slower rate, as the aggregate quantities slowly approach their stationary values. After 30 years, which is the expected working life of a household, all variables exhibit an elasticity of above 1, suggesting that the internal dynamics of the model have amplified the initial shock over time.

Moving on to the elasticities of the Gini coefficient, the initial impact of the shock is more pronounced, with an elasticity of -0.51 for human capital investment and -0.23 for consumption. This suggests that the inequality measured by Gini increases in response to the shock. The medium-term elasticities also remain strong, with a slight increase in the absolute magnitude of the response over a 5 and 10-year horizon. The elasticity of human capital inequality increases to -0.37 at the 10-year mark, and -0.43 after 20 years, indicating that the initial inequality shock has worked to increase human capital inequality over time. The elasticity of gross earnings remains low, at -0.06 after 10 years and only increases slightly to -0.1 after another 10 years. The response of consumption and assets to the shock is also negative and increasing over time, with the elasticity of assets reaching -0.4 after 20 years, while consumption inequality shows an elasticity of -0.29 after 20 years. However, overall the impact of the MIT shock on inequality remains less pronounced with most elasticities remaining below 0.5 in absolute value, in the long run.

Table 6: MIT Shock Elasticities

Mean	Impact	5 years	10 years	20 years	30 years	∞
HC Investment	0.385	1.056	1.313	1.448	1.487	1.523
Human Capital	0	0.15	0.404	0.797	1.024	1.293
Gross Earnings	1	1.3	1.578	1.921	2.105	2.318
Consumption	0.164	0.502	0.775	1.09	1.254	1.448
Assets	0	0.773	1.315	1.888	2.176	2.503
Gini	**Impact**	**5 years**	**10 years**	**20 years**	**30 years**	∞
HC Investment	-0.513	-0.598	-0.724	-0.911	-1.015	-1.169
Human Capital	0	-0.223	-0.372	-0.426	-0.431	-0.471
Gross Earnings	0	-0.037	-0.06	-0.104	-0.135	-0.179
Consumption	-0.228	-0.228	-0.25	-0.289	-0.311	-0.395
Assets	0	-0.195	-0.302	-0.403	-0.456	-0.601

Note: Elasticities in response to an unexpected, one-off 1 % TFP shock.

5.3 Comparison with a standard Aiyagari model

In order to further assess the role of human capital in propagating aggregate shocks, I compare the MIT shock response of the human capital model with the standard Aiyagari model that was first introduced in the previous section. Figure 8 compares the two models across the dimensions of assets, consumption and earnings. With respect to assets, the standard Aiyagari model predicts a smaller drop in average wealth than the human capital model and also suggests a quicker recovery. This is likely related to the additional loss of earnings in the human capital model, associated with lower human capital accumulation from the households. Consumption on the other hand falls much less in the augmented model, at least initially, which is likely a result of the lower level of consumption, increasing the willingness of households to draw down savings. Earnings fall by the same amount initially, but recovery is much faster in the standard model, as households do not need to replenish lost human capital. An interesting observation is that a uniform productivity shock can actually generate increases in earnings inequality in the human capital model, even though the effect is moderate at best.

Figure 8: Comparison of MIT shock response with standard Aiyagari model

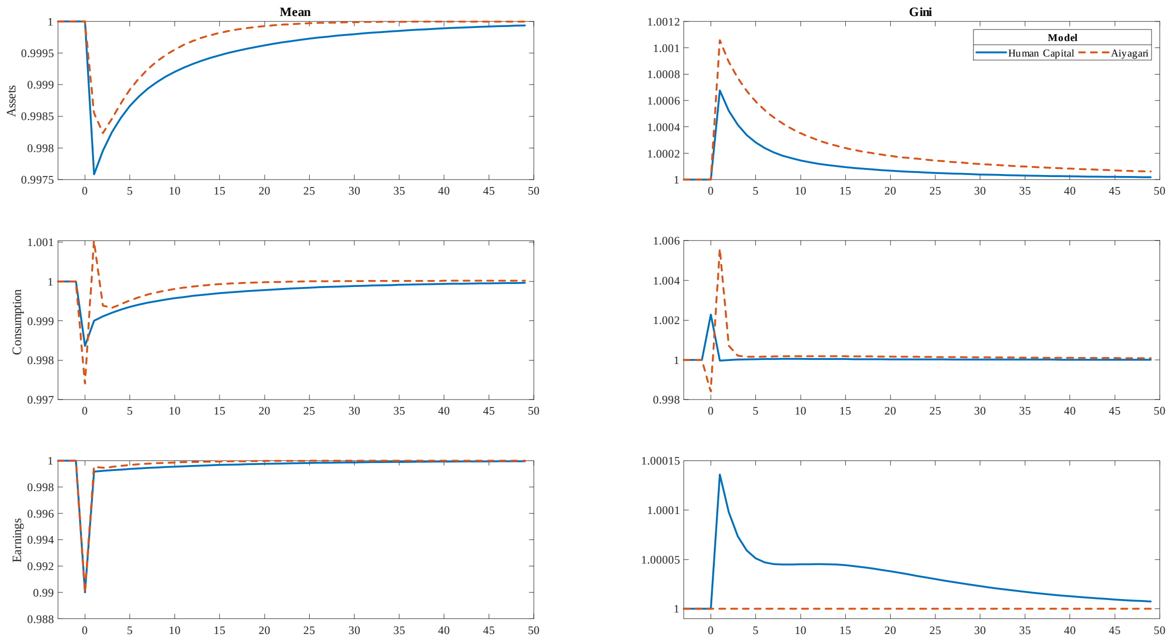

Note: All values are normalised to 1 in the stationary equilibrium.

5.4 Effects across the distribution

The last subsection explored the aggregate consequences of an unexpected economic shock. The analysis showed that even a perfectly symmetric shock can have long-lasting consequences for the distribution of human capital, wealth and consumption. As was explored in the previous section, households in this model behave differently, depending on their levels of wealth and human capital, and as a result, we should expect households to also react differently to the arrival of the shock. In this subsection, I explore how the impact of the MIT shock affects households at different parts of the wealth and human capital distribution.

Figure 8 plots the MIT shock response of human capital and gross earnings for households with different levels of pre-shock wealth and human capital. In spirit this reproduces Figure 4 above, with the modification that the reference group for each household is now a household with identical wealth, living in a world, where the shock never materializes. Specifically, the lines in the first row of Figure 8 correspond to the following sequences:

$$\frac{E_0(\{h_t\}_0^T|h_0^{pc},a_0^{pc},e_0) - E_0(\{h_t^{ss}\}_0^T|h_0^{pc},a_0^{pc},e_0)}{E_0(\{h_t^{ss}\}_0^T|h_0^{pc},a_0^{pc},e_0)}, \tag{2.18}$$

where the superscript ss indicates that the sequence is derived under steady state conditions. The other sequences in Figures 9 and 10 are calculated in the same manner, replacing h_t with the relevant endogenous variable in each case. For the interpretation of the findings of Figure 8, this means that the plotted sequences capture the heterogenous effect of the MIT shock on human capital and earnings of households. Given that considerable dynamic heterogeneity exists, even in the absence of a shock (see Figure 4 and analysis in the last section), the overall effect is likely to be exacerbated by the heterogeneous response.

We begin our analysis by looking at each graph in the first row of Figure 9 in turn. Each subplot reveals a smooth deviation from steady state human capital levels, irrespective of initial conditions. This pattern matches the one observed for aggregate human capital of the economy, with deviations reaching a trough of between 0.05 and 0.12 percent around 6 or 7 periods after the arrival of the shock. Moving from left to right, two observations become apparent: i) as we increase the initial human capital of households, the effect of the shock becomes less pronounced for all households, irrespective of initial wealth; ii) the difference in the impact on the lowest wealth households and the remainder becomes smaller. Both of these are a consequence of higher labour earnings, which help to relax the budget constraint for wealth-poor households.

Continuing on the last point, one of the most striking observations from each panel is the huge difference in the human capital response of households with the 20th percentile of wealth versus households with higher initial levels of savings. These wealth-poor households experience a dramatically bigger drop in human capital amounting to around twice that experienced by their more affluent peers. The effects are highly persistent for around 20 years, suggesting high overall cumulative losses. Moving up the distribution of human capital does reduce the absolute size of the drop for all households, but the relative gap remains. Echoing the analysis in the previous section, the nonlinearity in the MIT shock response seems to be highly concentrated at the left tail of the wealth distribution. Those with initial wealth equal to the 40th percentile, behave much more like households with wealth equal to the 60th or 80th percentiles, than their counterparts close to the borrowing constraint.

A loss of human capital will make itself known through lower pre-tax earnings. The second row of Figure 9 plots relative changes on gross earnings trajectories for the different households. Since the initial impact of the TFP shock dwarfs later more subtle dynamics, each plot contains an inset zooming in on the effects beginning one period after the shock. The results understandably mirror those of the human capital case. There is a small, but persistent drop in gross earnings relative to the counterfactual without a shock. Again, the biggest effects are concentrated on households with low initial wealth, who experience the biggest and most persistent deviations from their steady state trajectory. Increasing initial human capital does reduce the effect on all households, but fails to eliminate the gap.

Figure 9: Human capital and earnings response to MIT shock, by wealth quantiles

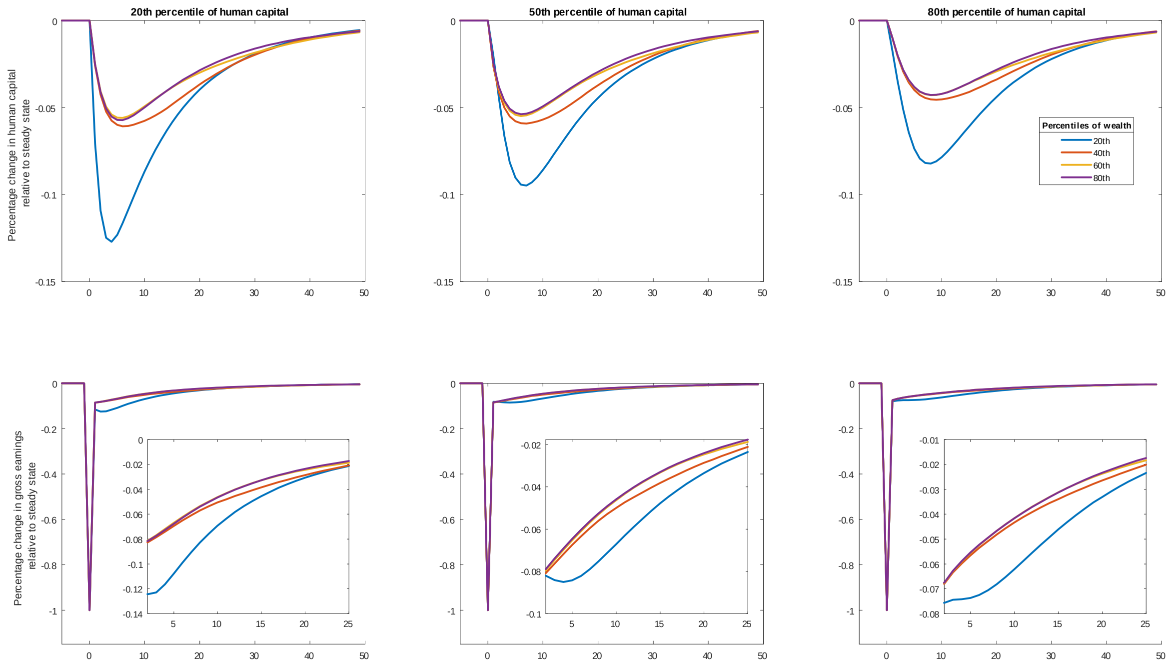

Note: Workers begin with mean level of productivity in the initial period.
Series normalised to expected progression in stationary equilibrium.

To analyse some of the reasons behind the strongly skewed response of human capital, Figure 10 plots the dynamics of consumption, human capital investment and assets for the same wealth and human capital quantiles as Figure 9. Looking at consumption in the first row, we see that initially the set of borrowing-constrained households actually appear less affected than their richer counterparts - or at least they do not appear considerably worse off. Their consumption schedule only falls behind of the rest of the households in the medium run as the lower earnings take their toll. Whilst appearing counterintuitive at first, the fact that lower-wealth households should be better able to insure their consumption against the shock is easily explained when considering the shape of the household's utility function: Lower-wealth households have - *ceteris paribus* - a lower level of consumption. Due to the concavity of $u(c)$, the same proportionate reduction in consumption leads to a larger loss in utility, and their incentive to maintain consumption is, therefore, stronger than for their wealthier counterparts.

To achieve this, and being unable to borrow additional funds, households with low wealth instead divert resources away from human capital investments into consumption, as the second row of Figure 10 shows. The poorest households experience a much more pronounced and much more persistent decline in their human capital investment levels than even those households with wealth equal to the 40th percentile of the steady state wealth distribution. This repurposing of investment resources, does provide some short-term relief, but leads to a loss of earnings in the medium run.

The same pattern is apparent when considering household wealth in the last row. The unexpected drop in incomes has a disproportionate effect on the poorest households, who draw down their small savings by an additional 7 to 5 percent relative to their steady state values. Households with richer initial endowments experience much smaller relative losses, meaning that they will be much better prepared to face future shocks to their incomes, be they idiosyncratic or aggregate in nature.

Figure 10: Consumption, human capital investment and assets, by wealth quantiles

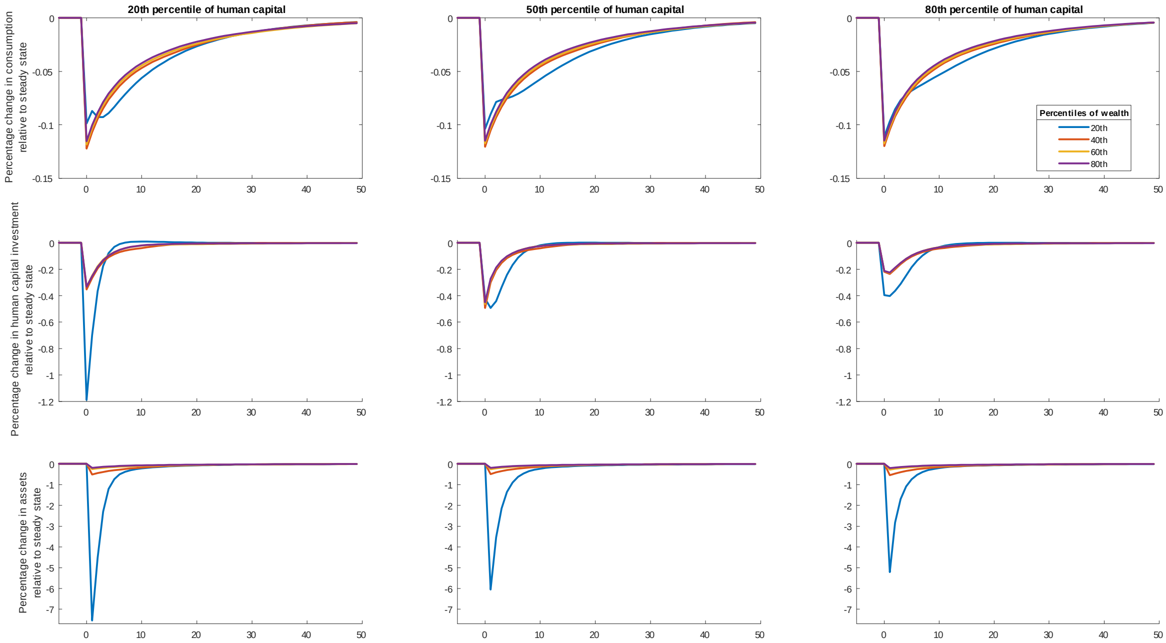

Note: Workers begin with mean level of productivity in the initial period.
Series normalised to expected progression in stationary equilibrium.

The analysis in this section has focussed on the effect of an unexpected, symmetric productivity shock on aggregate and distributional outcomes of the model. The analysis showed that even small shocks can have large consequences for aggregate variables of the model when assessed over the medium to long run. Further, it was shown that households close to the borrowing constraint have the strongest reaction to the shock, losing more human capital and earnings than even only moderately wealthy households, leading to persistent increases in human capital and earnings inequality with consequences for current and future consumption and welfare. This means that the initial distribution of wealth and human capital matters considerably for the aggregate response to an aggregate shock. A higher share of borrowing-constrained households implies a bigger, more persistent drop in human capital and thus a slower recovery.

All these conclusions are derived in a situation where the recession affects all households equally, irrespective of individual circumstances. However, economic shocks are hardly ever this blind. Often the effects of a recession fall disproportionately on those workers who are already struggling with their economic fortunes. In these cases, the impact of the shock is not distributed evenly but rather concentrated on a specific part of the pre-crisis productivity distribution. It is easy to extrapolate that under such circumstances the effect of a shock is going to be exacerbated, necessitating a policy response to protect the most vulnerable in society. The next section will address some of these points by exploring one of the biggest economic shocks in living memory which lead to unprecedented levels of public insurance - the Covid-19 pandemic.

6 Covid-19

The Covid-19 pandemic can justifiably be seen as the defining economic crisis of our age. At the very least, the pandemic caused a massive reduction in economic activity, largely as a result of lockdown measures put in place to protect public health. The corresponding loss of incomes and employment for millions of households necessitated unprecedented levels of intervention on the part of governments all across the world. As we slowly emerge from the pandemic, the question of how effective these policies were in avoiding catastrophe as well as the potential medium- and long-run impacts of the pandemic are at the forefront of a large research agenda.

In this last substantive section, I apply the model to these questions, by simulating the effect of the Covid-19 shock on the UK economy. In doing so I explore the effectiveness of the actual insurance policies applied by the government as well as an alternative, untargeted redistribution mechanism. I also use the model to assess the potential medium- and long-run impacts of the pandemic on human capital and wealth inequality. As I have shown in previous sections, the model is well suited to exploring the economy's response across these two dimensions, which are particularly relevant for the questions related to the economy's recovery from the pandemic. Given the recency of the pandemic, empirical research on the topic is still partially hampered by a data lag, and as a result, the impact of Covid-19 on human capital inequality in the UK is not yet fully understood. Initial research suggests that it has had a negative effect on education, and might have amplified human capital inequalities (Blundell et al. (2022)). However, any medium- and long-run impacts of these turbulent years will only be revealed with time.

In this situation, a model-based quantitative analysis can add value by providing a consistent story of how the economy might evolve over the coming years and help us to assess whether current policies are sufficient to address the long-term impacts of the pandemic, and to identify effective strategies for future crises. We have seen in previous sections, that the interaction between wealth and human capital inequalities, provides an amplification mechanism for aggregate shocks, which has the potential to deepen preexisting inequalities and slow down a potential economic recovery. Given both, the size of preexisting inequalities in the UK at the beginning of the Covid-19 pandemic, and the size of the shock, the effects are likely to be large and long-lasting. On the other hand, unprecedented policy intervention provided insurance to millions of workers and firms alike - a level of state-led welfare provision unmatched in any previous downturn. These measures were intended to stave off the worst effects of the pandemic and to ensure that the economy could bounce back quickly when conditions improved. Whilst ensuring consumption levels in a period of unprecedented economic turmoil was evidently in the interest of politicians and society at large, there remains the question of unintended consequences. A negative productivity shock creates disincentives to invest in human capital, *even if* incomes are perfectly insured. This means that it is possible, that even though the immediate effect of the pandemic caused a compression of the income distribution, the near future will see a rise in inequality along the lines of wealth, human capital, and income.

The model developed in this paper is uniquely suited to perform this analysis: it distinguishes between the level and productivity of human capital and can therefore account for the "price effect" of a negative, non-uniform productivity shock; and it also allows households to use their savings to invest in human capital and therefore takes account of the role of wealth inequality for the propagation of the Covid-19 shock. In this section I simulate the response of the model economy to an unexpected, non-uniform productivity shock, under different policy responses, providing valuable insight into the effectiveness of these policies.

6.1 Covid shock & policy responses

Covid-19, and the response intended to contain it lead to massive drops in economic activity, that reoccurred across much of 2020/21 and the spring of 2022. For 2020 alone the Bank of England estimated a drop in aggregate output of 9.7% (Harari & Keep (2021)), constituting a massive loss of earnings and employment for the average household. Generally, the effect was not uniform across the distribution of income. Much of the research on the economic impact of Covid-19, documents that households on lower incomes were hit harder than their contemporaries further up the income distribution (see Adams-Prassl et al. (2020), Blundell et al. (2020)). One main reason why initially poorer households were more affected by the pandemic is that they were more likely to work in industries that had been hit hardest by the pandemic, such as hospitality, tourism, and retail which experienced the closure of many businesses and the reduction in demand for services. As a result, these workers have been more likely to suffer from job losses and earnings reductions.

At the same time, richer households were more likely to be in industries that are less affected by the pandemic, such as finance and technology. As a result, their job security and earnings have been much more stable, particularly since remote working capabilities enabled many office-type professions to continue working throughout the pandemic. The overall picture that has emerged was one of a highly skewed income shock, that exacerbated existing inequalities along the lines of age, gender, health, education, professions and social class (see Marmot & Allen (2020), Marmot et al. (2021) and references therein).

Fortunately for many, a large part of the earnings losses were insured away by the government in an unprecedented public insurance effort, instituting the furlough scheme and business support initiatives across much of the pandemic years. Overall, this had the counterintuitive effect of the pandemic resulting in an actual fall in income inequality (see Blundell et al. (2022), Stancheva (2021)).

To accurately assess the effect of the Covid-19 pandemic the shock to the model economy will need to accurately reflect both the large, unequal reduction in pre-tax earnings and the generous income support schemes. My baseline calibration will take into account the large non-uniform drop in pre-tax earnings, as well as the relatively low reduction in after-tax incomes. Then I will calibrate two additional counterfactual insurance policies to assess the effectiveness of the government's response.

The starting point is going to be the numbers reported by Brewer & Tasseva (2021), who analyse the economic effect of the first wave of Covid-19 using data from the Family Resource Survey, the Understanding Society COVID-19 Study, and the UKMOD tax-benefit system. I use the numbers reported in Figure 1 in their paper. Specifically, I will use the drop in earnings (dark blue bar),[25] and net income after all taxes and transfers have been accounted for (black line with circles) as targets in my shock calibration procedure.[26]

To implement the shock across the distribution, I first assess which state (h_t, a_t, e_t) puts a household into which decile of the stationary (aka pre-Covid-19) earnings distribution. Then to calibrate the shock, recall that pre-tax earnings (q_t) are given by:

$$q_t = w_t e_t * \log(1 + h_t) \tag{2.19}$$

and post taxes and transfers income is:

$$y_t = \tau_1 q_t^{1-\tau_2}. \tag{2.20}$$

I define a new pandemic distribution of productivity e_t^{shock} and also a lump sum pandemic payment from the government $b_t^{shock} > 0$. As a result, during the pandemic households' pre- and post-policy earnings are calculated using modified versions of the previous equations: Gross labour earnings are calculated as

$$q_t^{shock} = w_t e_t^{shock} * \log(1 + h_t) \tag{2.21}$$

and post taxes and transfers labour income as follows:

$$y_t^{shock} = \tau_1 q_t^{shock^{1-\tau_2}} + b_t^{shock}. \tag{2.22}$$

Note that b_t^{shock} is a payment received in addition to the automatic tax and transfer stabilizers that are already embedded in the model. I assume that the government can finance its budget throughout the pandemic via additional borrowing without raising additional taxes - indeed this is broadly what happened in 2020/21.[27]

Both e_t^{shock} and b_t^{shock} are assumed to be functions of the exogenous productivity state e_t which effectively is a proxy for a worker's pre-pandemic occupation. As was noted above, the impact of Covid-19 was highly skewed towards workers in low-income (productivity) occupations, making pre-pandemic productivity a relevant predictor of how impacted a household would be by the pandemic.

[25] I exclude self-employed earnings, in line with my sample selection.

[26] Figure 1 in Brewer & Tasseva (2021), shows average monthly estimates. Since they focus on April and May of 2020 that means that extrapolating to an annual frequency might slightly overestimate the impact of the pandemic.

[27] The aim of this paper is to study the impact of Covid-19 and the government's response in the short to medium run. For simplicity, I ignore the implications of this increase in government debt in the remainder of this analysis. However, I acknowledge the importance of this question and provide some discussion of the limitations that this brings in the later parts of this section.

I approximate both e_t^{shock} and b_t^{shock} using polynomial functions, and try and minimize the squared percentage difference between the reported drops and those implied by the model, taking into account the general equilibrium effect of changes in the wage rate.[28]

Figure 11 shows the model-implied changes in pre- and post-policy incomes for the calibrated values of the pandemic productivity shock and the emergency subsidies and compares these changes with those reported by Brewer & Tassseva (2021). The fit of the calibration is quite good across the distribution of pre-pandemic gross earnings, reproducing a slightly skewed "U" shape in gross earnings losses and a monotonously decreasing, almost linear shape for net earnings.

Figure 11: Pre- and post-taxes and transfers earning drops

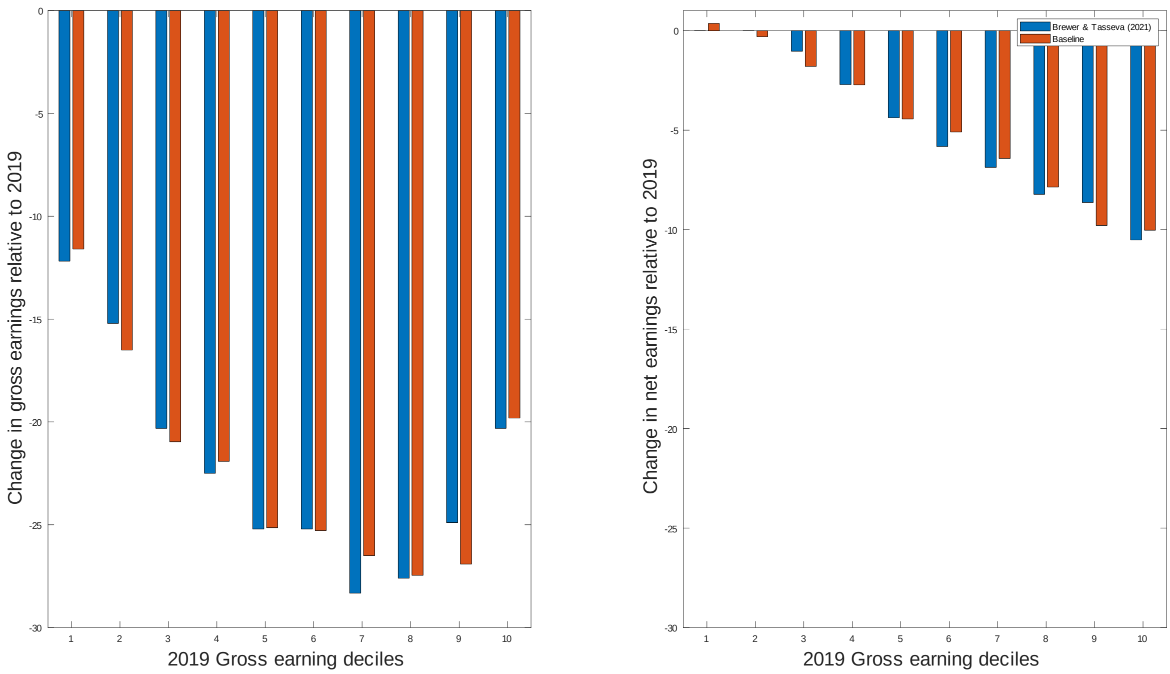

Note: Data points transcribed from Brewer & Tasseva (2021), Figure 1.

To complement the baseline calibration, I devise two additional policy scenarios: i) no intervention and support measures and ii) a uniform lump-sum payment. The no-intervention scenario is a crucial counterfactual to evaluate the effectiveness of the observed intervention in 2020/21. For this purpose, I assume that the government did not provide any additional Covid-19 relief to households, but simply relied on the automatic stabilizers to offset the loss in employment and earnings. For this, I set $b_t^{shock} = 0$.

Emergency support measures linked to the Covid-19 pandemic were heavily targeted

[28] See Appendix for a detailed description of the calibration.

and income-contingent. The Coronavirus Job Retention Scheme ("furlough") for example provided employers with government grants in order to retain workers at 80% of their usual wages (albeit subject to a maximum payment). This policy clearly benefitted high-wage workers, as the absolute subsidy they received would have been considerably larger than for low and medium-wage peers. Other interventions, such as the £20/week universal credit uplift were targeted at the lower end of the income distribution. On balance, these interventions were designed to provide relief to those that required it most. I counterpoint this by assuming that an agnostic non-contingent benefits policy had been put in place instead, which just provided every household with the same amount of relief irrespective of circumstances. To implement this, I calibrate the emergency benefit, so that the total amount that was spent on b_t^{shock} in the baseline case gets distributed equally across households. The corresponding net income changes for both policies can be found in Figure A1 in the Appendix.

One of the most noticeable impacts of the Covid-19 pandemic were restrictions on the way that households could spend their money. Frequent lockdowns meant that nonessential shops were closed for large parts of the year, socialising in bars, cafes and restaurants was heavily curtailed leading to large reductions in consumer spending. Evidence for the UK and other economies imposing lockdown restrictions suggest, that these policies lead to large reductions in household consumption, particularly for higher-income households (see for example Hacioglu-Hoke et al.(2021); Bank of England (2020); Tenreyro (2021) for the UK; Dossche and Zlatanos (2020) for the EU; and Miescu and Rossi (2021) for the US). Specifically, Davenport et al. (2020) found that among the two highest income quintiles, consumption decreased by about 25% in the early months of the crisis, with smaller changes for lower income groups (consistent with patterns reported in Bank of England (2020) and Tenreyro (2021)). In the context of the model such forced savings are likely to lead to a reallocation by richer households who are constrained in their consumer spending towards financial savings (see on this point Angelopoulos et al. (2021)) and human capital investments. In order to capture this important mechanism, I add a consumption ceiling to the model which applies during the initial phase of the pandemic.[29]

[29] The consumption ceiling imposes an additional utility cost on consumption above a certain threshold. I calibrate both the consumption ceiling and the penalty to match as closely as possible the relative mean decrease in consumption for the five pre-Covid-19 income quintiles as reported in Davenport et al. (2020). For details see Figure A7 in the Appendix. For the computational algorithm, see Appendix.

6.2 Impact of Covid-19 on aggregate economic variables

I take the calibrated shocks and apply them fully to 2020, and then again half of the productivity shock and half the extra benefit payment to 2021.

I first discuss the medium-run predictions for the aggregate economy under the baseline calibration and compare them to the no insurance and uniform insurance counterfactuals. After this, I analyse the differences in the response of the aggregate economy through the lens of the heterogenous responses of households at different points of the pre-Covid-19 wealth and human capital distributions. The exercise shows that the initial state of households plays an important role in determining the households' response to the Covid-19 pandemic, and interacts with the policy to drive the dynamics of the aggregate economy.

The Covid-19 pandemic had large and immediate consequences on the UK economy. Many sectors and industries were partially, or fully locked down, leading to the closures of thousands of businesses and forcing millions of laid-off workers to rely on state-provided benefits. But while the prospect of future full-scale lockdowns appears unlikely, there remains some uncertainty of when, or if, the economy will truly recover. Large economic downturns can leave lasting scars on an economy, by destroying burgeoning businesses or causing mass layoffs leading to the destruction of match-specific human capital (see for example Ouyang (2009), Huckfeldt (2022)). Job retention programmes such as the furlough scheme, were designed to maintain employment relations and avoid the painful effects of job separations on a large scale. Yet still, a loss of "business as usual" for large parts of the years will have certainly affected the opportunities and incentives for learning, training and skill acquisition, which are crucial in building human capital.[30]

Figure 12, plots the response of the main endogenous outcomes of the model in response to the Covid-19 shock using the baseline calibration. I focus primarily on investment into and the stock of human capital since this is the mechanism through which the pandemic might have long-run effects on future economic prosperity. On impact investment into human capital falls sharply, reaching a low of -4% relative to pre-Covid-19 levels. Recovery to steady state levels takes around 10 to 15 years, even though most of the initial investment level is reclaimed by 2030. Average human capital on the other hand falls much more gradually, reaching a trough of -0.5% in the second half of the 2020s. The recovery is even more gradual, and human capital levels are projected to be around 0.3% below 2019 levels until at least 2050. This hints at the likelihood of a very long-winded

[30] Admittedly, for many of us the pandemic necessitated learning many new skills, such as remote working, teaching and collaborating. Whether these adaptations merely helped us buffer the worst excesses of the recessions, or will make us more productive in the future is an interesting and open question.

recovery, with economic potential below pre-crisis levels for a generation.[31]

On a slightly more positive note, it appears that the projected path of the economy points towards greater levels of equality - at least in the medium run. The gini of human capital investment spending drops sharply in 2020, by almost 5 pecent. While investment inequality recovers fairly quickly, the distribution of the stock of human capital becomes more equal throughout the mid-to late 2020s. Reaching a trough of about 0.2% below steady state gini levels, human capital inequality only recovers gradually, staying below pre-pandemic levels for around 10 years, reaching its steady state value around 2035. Interestingly, human capital inequality does not return simply to its steady state level but continues to grow steadily into the second half of the century. This is likely driven by the large and sustained increase in wealth inequality. While mean assets drop by around 2.5%, wealth inequality increases by around 1.5% and remains elevated for over 25 years (see also Angelopoulos et al. (2021) for evidence of increased wealth inequality post-Covid-19). We have seen in earlier sections of this paper that wealth inequality can lead to persistent increases in human capital inequality, so it is likely that this increase in asset inequality is driving the steady increase in human capital inequality noted above. Average consumption and consumption inequality drop dramatically in 2020 as a result of the imposed consumption restrictions but recover relatively quickly to their pre-Covid-19 levels after that.

[31] The model is based on the human capital of working-age individuals and therefore cannot speak to the skills and abilities of future cohorts of labour market entrants. However, given the large-scale disruptions to normal teaching schedules at all levels, the projections of the model are likely on the optimistic side.

Figure 12: Response of aggregate variables to Covid-19, Baseline

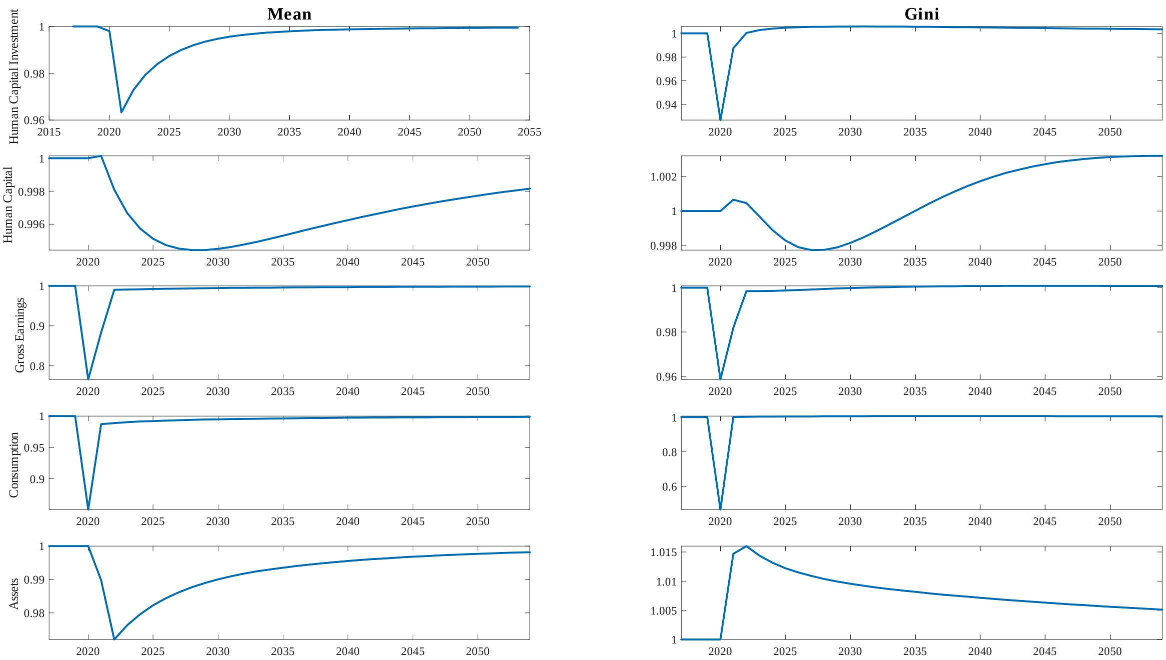

Note: Baseline Covid-19 shock calibration.

Series normalised to expected progression in stationary equilibrium.

At the aggregate level, the model predicts a loss of human capital in the medium run, leading to a loss in economic output and growth potential. Consumption restrictions lead to more equal consumption - at least for a brief period - but the resulting response of high-income households leads to increased wealth inequality which feeds an increase in human capital inequality in the long run.[32]

But how much of this is due to the nature of the economic shock, and the way it has affected workers across the distribution of incomes differently, and which part is played by the government's targeted Covid-19 relief policies? To assess the effectiveness of targeted intervention to help cushion the blow of the Covid-19 shock, and distribute the associated economic pain more equitably, Figure 13 reproduces Figure 12, adding the no insurance and uniform insurance counterfactuals.

Under the no insurance scenario, the government did not provide any additional payments to households but merely allowed the automatic stabilizers of the tax and transfer system to operate as in normal times. As a result, the economic impact of the pandemic

[32]Indeed the response of the asset gini is around 3 times larger in the model with a consumption cap compared to a comparable model without such a restriction. See Figure A8 in the Appendix for details.

is much bigger for all households, but additionally, the removal of the broadly regressive Covid-19 benefits leads to a bigger pass-through of the shock to lower-income households. The proposed policy of providing benefits uniformly across the distribution of households, has very similar dynamics relative to the baseline calibration (see Figure 13, Table 7, and Figure A2 in the Appendix), whilst potentially being easier to administer. Since the uniform insurance scenario is qualitatively and quantitatively similar to the baseline case for the aggregate behaviour of the economy, I will focus on the comparison of the baseline case and the no insurance scenario for now.

Comparing the two scenarios, it is evident that the impact of the Covid-19 pandemic would have been a lot more severe in the absence of additional policy measures. The projected impacts on human capital, and assets are both more negative and much more persistent in the absence of emergency support measures. More importantly, the slightly equality-enhancing tendencies that were documented for the baseline case, reverse in sign, with inequality increasing particularly for human capital and wealth over the short and medium run. This qualitative difference is very suggestive of the success of the government's targeted approach in avoiding a large potential increase in inequality after the pandemic.

Figure 13: Response of aggregate variables to Covid-19, alternative policies

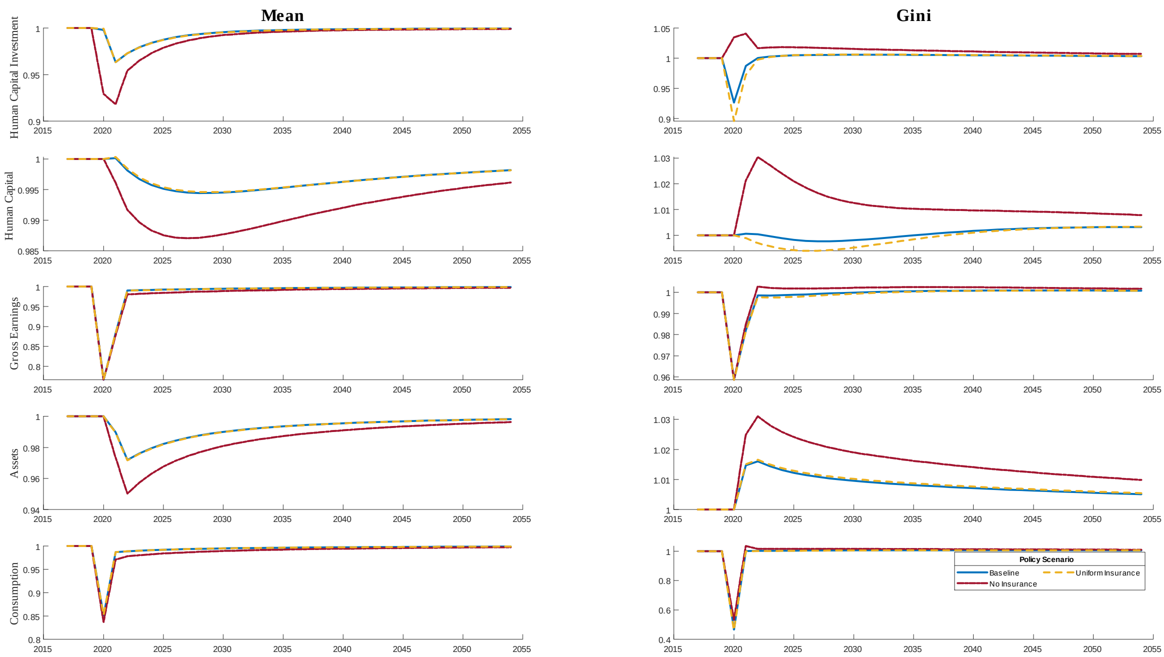

Note: Baseline, uniform and no insurance Covid-19 shock calibration.
Series normalised to expected progression in stationary equilibrium.

As I have illustrated with the analysis of Figures 12 & 13 above, many of the negative consequences of the Covid-19 pandemic are highly persistent. Given the slow return to steady state conditions, simply documenting the peaks and troughs of the endogenous variables does not provide a complete picture of the true impact of the crisis. In order to assess the cumulative effects, Table 7 provides a selection of cumulative effects of the Covid-19 shock across different time horizons and policy scenarios. The values presented in Table 7 are expressed as cumulative percentage deviations of the target outcome relative to an annual value of that outcome in the pre-Covid-19 stationary distribution. Specifically, the cumulative effect of the Covid-19 pandemic on outcome z at time horizon T is calculated as:

$$C_z^T = 100\% \sum_{t=0}^{T} \frac{z_t - z_{ss}}{z_{ss}}, \tag{2.23}$$

In the case of slow-moving variables, such as the aggregate stock of human capital, the impact of the Covid-19 pandemic will show itself over the course of many years and even decades. This means that the impact might be less detectable, as the effect is spread out across many years, but that does not mean the total effect is necessarily negligible. We know from the MIT shock experiments conducted in the previous section, that human capital responds slowly to the shock, but carries a significant degree of inertia. This property is transferred to the case of the Covid-19 shock. In 2020, human capital investment remains roughly constant given the baseline calibration, while after 5 years the reduction is equal to 10% of annual human capital investment pre-Covid-19. This is likely due to the fact that in 2020, consumption restrictions incentivise households to channel resources into skill accumulation rather than consumption, balancing out the shortfall in the short run. Aggregate human capital responds slowly to this fall - reaching a total reduction of 1% after 5 years which then accelerates to -4% after 10 years. In the long-run this cumulative loss of human capital is far from trivial. After 30 years the economy has lost human capital worth 11% of the aggregate capital stock of the pre-Covid-19 economy amounting to over $\frac{1}{3}$ of a percent per annum.[33]

The lesson to be drawn from this analysis is that much of the adjustment of human capital lies in the future. Even if many other indicators of economic activity appear to be converging back to normal, we should be wary of the long-term implications of this slow-moving variable, as any further shocks to the economy can compound the loss of human capital and have lasting effects on economic outcomes. It is therefore important to plan ahead and implement policies that can mitigate the potential long-term effects of

[33] A baseline calibration of the model without the consumption restriction predicts losses almost exactly twice that order, suggesting that the forced savings in the short run pay off in terms of much smaller losses to aggregate human capital in the long run. See also Figure A8 for a comparison.

the Covid-19 pandemic on human capital, to ensure that the economy can recover and thrive in the years to come.

Trying to reduce the immediate impact of the pandemic was the main aim of the Covid-19 support policies, but as the comparison of the baseline calibration to the no insurance counterfactual suggests, it also helped alleviate some of the negative long-run effects. In the short- and medium-run, human capital investment drops 7% on impact and a cumulative 26% over the first 5 years, in the absence of additional Covid-19 support measures. This additional reduction in investment into human capital leads to a much larger and more persistent drop in the aggregate stock of human capital, falling 3% after 5 years, and 10% after 10 years. A generation after the initial onslaught of the pandemic, the economy would have lost an *additional* 16% of the annual pre-pandemic human capital stock, in the absence of the additional Covid-19 benefit payments.

Table 7: Cumulative effects of the Covid-19 shock

Baseline

Mean	Impact	5 years	10 years	20 years	30 years	∞	Gini	Impact	5 years	10 years	20 years	30 years	∞
HC Investment	0	-10	-14	-17	-18	-20	HC Investment	-7	-8	-5	0	5	21
Human Capital	0	-1	-4	-8	-11	-17	Human Capital	0	0	-1	-1	1	21
Gross Earnings	-23	-38	-41	-45	-48	-51	Gross Earnings	-4	-6	-7	-6	-6	-2
Consumption	-15	-19	-23	-26	-28	-32	Consumption	-53	-53	-51	-45	-40	-12
Assets	0	-8	-15	-22	-26	-29	Assets	0	6	11	20	26	56

No Insurance

Mean	Impact	5 years	10 years	20 years	30 years	∞	Gini	Impact	5 years	10 years	20 years	30 years	∞
HC Investment	-7	-26	-33	-37	-39	-43	HC Investment	3	13	21	35	45	77
Human Capital	0	-3	-10	-20	-27	-37	Human Capital	0	10	19	29	39	79
Gross Earnings	-23	-41	-48	-56	-61	-68	Gross Earnings	-4	-5	-4	-2	0	7
Consumption	-16	-25	-32	-40	-44	-51	Consumption	-47	-39	-31	-17	-5	47
Assets	0	-16	-29	-42	-49	-57	Assets	0	11	22	39	51	106

Uniform Insurance

Mean	Impact	5 years	10 years	20 years	30 years	∞	Gini	Impact	5 years	10 years	20 years	30 years	∞
HC Investment	0	-10	-14	-17	-18	-20	HC Investment	-10	-13	-10	-4	1	19
Human Capital	0	-1	-3	-8	-11	-17	Human Capital	0	-1	-4	-6	-4	18
Gross Earnings	-23	-38	-41	-45	-47	-51	Gross Earnings	-4	-7	-7	-7	-7	-3
Consumption	-15	-19	-22	-26	-28	-31	Consumption	-54	-54	-52	-47	-41	-11
Assets	0	-8	-15	-23	-26	-30	Assets	0	6	12	21	28	59

Note: Cumulative differences in key endogenous outcomes due to the Covid-19 pandemic under different policy scenarios.
All numbers expressed as percentage differences relative to 2019 baseline.

For the dynamics of human capital in the UK economy following the Covid-19 pandemic and therefore the potential for future economic growth and development, the preceding analysis provides two key insights: First, while many other economic indicators, such as earnings and consumption begin to recover immediately after the pandemic shock subsides, human capital is much more sluggish, and thus the biggest effects in that regard might still lie ahead of us. The prospect of aggregate human capital staying one percent below its pre-Covid-19 level for around 20 years does not bode well for the future dynamism of the UK economy and its ability to adapt successfully to future challenges. On the other hand, while the outlook is somewhat bleak, decisive policy action has likely delivered us from an even worse fate. The policy counterfactual suggests that in the absence of additional Covid-19 support measures, the fall in aggregate human capital would have been much larger, and crucially coupled with an increase in human capital inequality, exacerbating pre-existing divides in society.

In the next section, I focus in on how households across the distribution of wealth and human capital react to the pandemic, in order to shine a light on some of the aggregate dynamics we have seen in this section, but also to assess where policy has been successful if it has, and what lessons can be learned for the future.

6.3 Impact of Covid-19 across the distribution

There are two main channels through which the Covid-19 pandemic interacts with preexisting inequalities: i) the productivity shock is unequally distributed across the distribution of pre-Covid-19 labour productivities; and ii) households with different initial levels of wealth and human capital have different capacities to deal with a surprise shock, even if it is symmetric in its impact. In this subsection I explore the role of initial conditions in response to the Covid-19 shock, providing some insight on the individual household behaviours that constitute the aggregate response of the economy that we discussed in the last subsection.

Figure 14 presents the impact of the Covid-19 shock on the human capital levels of households with different initial levels of wealth, human capital and productivity in the period preceding the shock. The figure is analogous to Figures 8 & 9 in the preceding section with two modifications to allow a better understanding of the complex interactions between the nonlinear shock, the policy response and wealth and human capital inequality. Firstly each row in the following figures refers to a policy scenario, so a comparison across rows will allow for an assessment of differences due to different policy responses, allowing a counterfactual evaluation of a policy's effectiveness. Secondly, in order to account for the nonlinear impact of the shock across different pre-Covid-19 productivities, each column represents a fixed level of the pre-Covid-19 productivity distribution. Comparison across columns therefore allows an assessment of the impact by pre-Covid-19 productivity. To account for the correlation between human capital and productivity, I fix the initial level of human capital to the mean level conditional on pre-Covid-19 labour productivity and wealth quantiles. Table 8 provides some summary statistics describing relevant features of the households conditional on e_t percentiles.

Table 8: Summary statistics conditional of productivity percentiles

pc(e_{2019})	$\bar{h}_{2019}$	pc($\bar{h}_{2019}$)	$\bar{q}_{2019}$	pc($\bar{q}_{2019}$)	Q1	Q2	Q3	Q4	Q5
10th percentile	1.435	24	0.494	11	0.294	0.191	0.194	0.166	0.154
50th percentile	1.556	43	0.886	49	0.182	0.203	0.222	0.2	0.193
90th percentile	1.752	69	1.624	88	0.069	0.189	0.245	0.243	0.254

Note: Summary statistics conditional of productivity percentiles pre-Covid-19.
A bar refers to the average value within the productivity percentile (pc = percentile).
"Qx" refers to the population shares that fall into the respective wealth quantiles.

Through the model, the endogenous variables - human capital, earnings and wealth are highly correlated with exogenous human capital productivity. Workers within the 10th percentile of productivity have on average human capital that would put them into the 24th percentile of human capital, compared to the 43rd percentile for median productivity and 69th percentile for the top 10% of productivity. In terms of gross earnings, these differences are even more striking. The productivity percentiles chosen correspond almost perfectly to the 10th, 50th and 90th percentiles of the gross earnings distribution on average allowing a neat comparison with the earnings losses documented in Figure 9. In terms of wealth, there are also clear disparities between productivity percentiles, with lower productivity percentiles exhibiting a larger share of households in the first two quantiles of the wealth distribution, with the distribution reversing as we consider the highest productivity percentiles.

Focussing on the baseline calibration, moving from left to right the first thing to notice is that the negative impact of the Covid-19 shock on the evolution of human capital is fairly similar for households from the 40th percentile of the pre-Covid-19 wealth distribution upward. Households beginning with assets equal to the 40th percentile of the pre-Covid-19 wealth distribution, reduce their human capital by a maximum of around 0.5% to 0.6% relative to their steady state trajectory by 2030 and then begin a slow recovery. The richer 60th and 80th percentile households lose slightly less, reaching a through of around -0.4% relative to their pre-Covid-19 benchmark irrespective of productivity percentile. The average effect - depicted by the black broken line in each subplot - is slightly progressive in the sense that across higher initial productivity states and levels of human capital, the losses for the average household become incrementally more negative.

Big differences are noticeable when we focus on households close to the borrowing constraint. Across most of the distribution of initial productivity and human capital, households beginning with assets equal to the 20th percentile behave distinctly from their richer counterparts - an expression of the nonlinear behaviour that was encountered in the case of the MIT shock as well.

At the lower end of productivity, the poorest households are the only group not to increase their human capital slightly relative to their steady state counterfactual. However in the medium run, their trajectory seems to be slightly advantageous relative to richer households, but the effect is marginal and likely driven by strong emergency benefits - a view that is confirmed when one compares the trajectory in the absence of additional benefits. For median and higher productivities, these households lose considerably more human capital compared to their pre-Covid-19 trajectory than their richer peers. Their losses amount to about twice what even moderately wealthy households lose, suggesting that the borrowing constraint significantly affects their ability to respond to the pandemic. At the higher end of the productivity distribution particularly, the relative change in human capital becomes strongly negative, exceeding -1% indicating that at the higher end of the earnings distribution borrowing-constrained households come off considerably worse than the rest. This is unsurprising, as at this point in the earnings distribution the net income effects of Covid-19 are at their most severe, implying that the self-insurance mechanism via personal and household savings plays a larger role.

In order to interpret these findings it is helpful to remind ourselves that what Figure 14 implies is not that wealth-poorer households accumulate more human capital than asset-rich households. Rather, the figures depict deviations from steady state trajectories (see Figure 4) for which we know that - *ceteris paribus* - richer households accumulate human capital at much faster rates than poorer ones. The broad patterns observed in the subplots of the first row of Figure 14, therefore hint at a speed-up of wealth-based inequality in human capital accumulation, even though there also appears to be a slight pattern of contraction of human capital accumulation at the top end of the productivity distribution. This could explain the pattern that was observed with respect to the aggregate development of human capital inequality in Figure 13. Initially, there is a short spike in human capital inequality driven by a reallocation of resources from consumption to human capital investment by very wealthy households. This is followed in the short run, by a contraction of the human capital distribution driven by a relative fall in human capital accumulation by households in higher-productivity occupations. Finally, in the medium run, increases in wealth inequality combined with a strengthening of the correlation between human capital and wealth, lead to a persistent increase in human capital inequality in the long run.

The second row presentes the evolution of human capital assuming no additional Covid-19 benefits. Immediately it is obvious that as a result the decline in human capital relative to the steady state is considerably larger, on the order of an additional 1% to 1.5% across the distribution. This additional loss of human capital is reflected on the aggregate level in the big decline in aggregate human capital that was observed in this case in Figure 13. As expected, the absence of additional benefit payments does affect borrowing-constrained households most severely. Generally speaking, wealth-poorer households experience larger relative declines in human capital than those with more assets, with borrowing-constrained workers experiencing particularly large relative falls. This tendency will on balance increase the correlation between wealth and human capital in the medium run. Furthermore, if we take account of the composition of households in each productivity percentile - e.g. much more low-wealth households in low productivity percentiles, and a higher share of high-wealth households in high productivity percentiles - this also accounts for the large raise in human capital inequality that was predicted in this case. An additional observation in this case, is that the average effect is approximately "U" shaped across the productivity distribution, mirroring the impact of the Covid-19 productivity shock.

Under the uniform insurance policy, there is a strongly progressive redistributive element, as low-income households benefit disproportionately from the uniform payment (see Figure A1 in Appendix). The effect of this policy can be seen in the third row of Figure 14. For households in the lowest productivity percentiles, there is a strong positive effect for households at the 20th percentile of wealth, compared to the baseline policy. These households benefit particularly from the additional resources provided by the Covid-19 benefits. Interestingly richer households do not appear to change their behaviour relative to the baseline calibration. These households are likely already unconstrained as a result of their savings and thus the additional resources from the policy do not significantly alter their human capital accumulation decisions. At the median level of productivity, the uniform policy is actually slightly worse in terms of post taxes and transfers earnings and as a result, wealth-poorer households reduce their human capital slightly more relative to the baseline policy. Finally, at the highest levels of productivity, net labour earnings are fairly similar between both policies, and as a result, the projected human capital trajectories are very similar.

Figure 14: Human capital response to Covid-19, by wealth quantiles

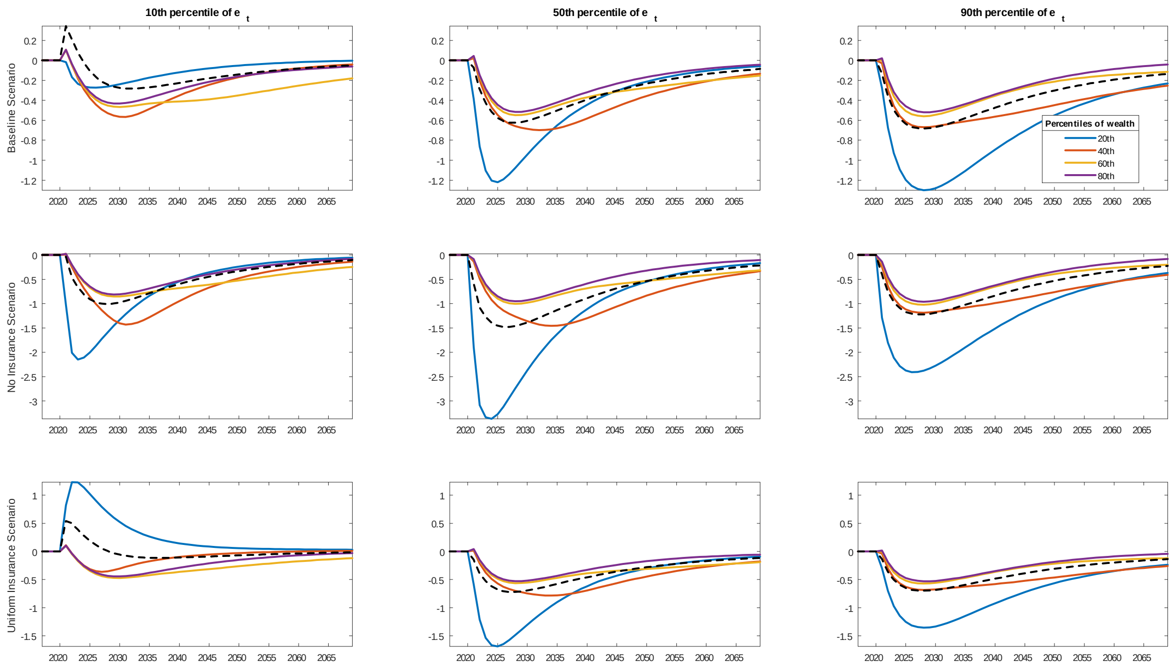

Note: Series normalised to expected progression in stationary equilibrium. Workers begin with human capital equal to the mean level of human capital within their wealth-productivity percentile. See Table 8 for details.

The analysis in this subsection has shown how the aggregate dynamics of human capital post-Covid-19 arise from the differential response of households across the distribution of productivity, human capital and wealth. Under the baseline scenario, I find that the response to Covid-19 is progressive in the sense that on balance households with lower levels of human capital, who begin in lowproductivity occupations and industries reduce their human capital less than households in higher productivity states who also have higher levels of human capital. This provides a compression of the human capital distribution in the medium run, leading to lower skill inequality. In the long run, however, a strengthening of the correlation between wealth and human capital leads to wealthier households recovering much quicker from the pandemic shock, using their substantial resources to rebuild lost skills, leading to an increase in human capital inequality that proves highly persistent.

Wealth, in this case, provides an important dimension of heterogeneity, as I find that the impact on households human capital accumulation depends strongly on their initial asset holdings, even when conditioning on the same initial productivity levels. Particularly low-wealth households seem to be vulnerable to the Covid-19 shock, particularly when they are not receiving additional benefit payments. Higher-wealth households on the other hand appear less sensitive, likely since they have adequate personal resources to respond to the crisis.

Figure 15 highlights the differential response in asset accumulation across different levels of initial productivity at the onset of the pandemic, by reproducing Figure 14, with a focus on households' wealth positions rather than human capital. Beginning with the baseline case in the top row, the most prominent feature is a disproportionate reduction in the asset positions of households with low wealth. While all households lose some wealth relative to their pre-Covid-19 trajectory, households close to the borrowing constraints lose up to over 70%, leaving them particularly vulnerable to future shocks. The effect is particularly strong for households in median or higher productivity states and proves fairly persistent - likely driven by lower human capital going forward.

Richer households experience much lower relative reductions, which are likely driven by two interrelated factors: i) households with some savings were better able to protect their stock of human capital and therefore their future earnings potential (see also Figure A4 in the Appendix), and ii) high wealth households were forced by consumption restrictions to divert some of their usual expenditures into human capital and financial savings (see also Figure A3 & A4 in the Appendix).

These effects are exacerbated through the absence of additional Covid-19 support measures as can be seen in the second row of Figure 15. Savings of low-wealth households are all but wiped out by the pandemic, while additionally there are also large relative reductions in the asset positions of moderately wealthy households with wealth around the 40th percentile.

The uniform insurance counterfactual is comparable to the baseline case, even though strong redistribution towards the bottom of the earnings distribution means that borrowing-constrained households at the bottom of the productivity distribution benefit somewhat from the transfer of additional resources.

Figure 15: Wealth response to Covid-19, by wealth quantiles

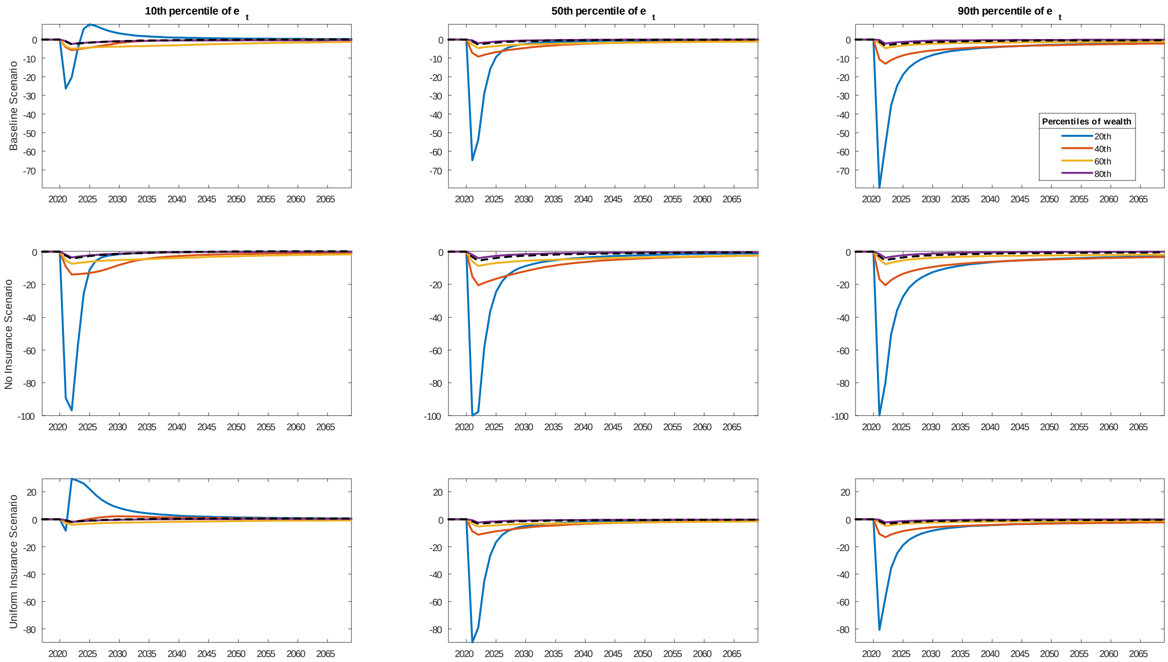

Note: Series normalised to expected progression in stationary equilibrium. Workers begin with human capital equal to the mean level of human capital within their wealth-productivity percentile. See Table 8 for details.

The changes in the relative asset positions of households are strongly inequality enhancing, as was also initially observed in Figure 13. For the medium and long-run dynamics of the economy this has two important implications: First, increased wealth inequality sows the seeds of increases in human capital and therefore earnings inequality in the future. The analysis of previous sections has shown that wealth-poorer households accumulate human capital at much slower rates than richer households, meaning that an increase in wealth inequality will lead to a slower economic recovery. Second, as low-wealth households fail to accumulate substantial buffer stock savings throughout the pandemic and post-pandemic period, they are not only in a weaker position when it comes to taking advantage of future opportunities, such as high-productivity jobs, but they are also much more vulnerable to future economic shocks.

The importance of self-insurance is particularly highlighted when we consider the counterfactual without additional Covid-19 benefit payments. If households were to solely rely on the automatic stabilizers, households with low wealth perform distinctly worse compared to the baseline case. Whilst households all across the distribution of assets do worse under this scenario, it is particularly borrowing-constrained households who are forced to reduce their desired human capital plans in response to the shock, ultimately leading to higher human capital inequality and also a stronger association of wealth and human capital. This suggests that while the Covid-19 benefits were designed to benefit low-income households, they particularly benefitted low-wealth households. This by itself raises some interesting questions about the design of future policies in response to aggregate shocks, and also the dismal topic of paying for the associated costs. I will discuss these questions and some other points in the next subsection.

6.4 Welfare effects and policy implications

The Covid-19 pandemic was much more than simply an economic shock. The loss of countless lives, the disruption of normal social relations for the majority of two years, and a general sense of uncertainty and anxiety have imposed psychological costs far in excess of any measurable loss of income, employment or consumption. Yet it is an economist's job to try his best and quantify the impact of the pandemic so that society might better understand what was lost, and how to decide what to do in the future.

In this section I will use the model predicted transitions post-Covid-19, to calculate the distribution of welfare losses over the course of the pandemic and its aftermath. For this purpose, I obtain the value functions at the beginning of the pandemic period. Since the value function encodes the solution to the infinite horizon problem of a forward-looking household, the value function at the beginning of the pandemic period summarizes not only the immediate welfare impact of Covid-19 but also the experience of the household along the transition path.

The value function is known at every point in the state space, so I can obtain the distribution of welfare across the dimensions of wealth, human capital and exogenous productivity. To make the welfare measure more interpretable, I convert the welfare differences into percentage differences of *lifetime equivalent consumption* in the stationary equilibrium, using the homogeneity of the utility function:

$$\lambda(a_t, h_t, e_t) = \left[\frac{V_0(a_t, h_t, e_t)}{V_{ss}(a_t, h_t, e_t)}\right]^{\frac{1}{1-\sigma}} - 1, \tag{2.24}$$

where $\lambda(a_t, h_t, e_t)$ is a number equivalent to the percentage of additional consumption that an agent in state (a_t, h_t, e_t) would require in every state to be indifferent between going through the Covid-19 crisis and post-pandemic period and staying in the stationary world without Covid-19 forever.[34] V_0 and V_{ss} refer to the value functions at the beginning of the pandemic and in the stationary equilibrium respectively.

Figure 16 shows the welfare losses due to the Covid-19 pandemic, both on aggregate and across the distributions of pre-pandemic earnings and wealth holdings. Overall the pandemic had a significant and negative effect on the welfare of households. Under the baseline scenario, the average household would have been happy to trade in a permanently lower level of consumption by 1.15% to avoid the pandemic. Under a uniform benefits policy, the number is marginally smaller, implying that on average a uniform policy would have been preferable. Indisputable however is that additional policy intervention is preferable to simply relying on automatic stabilizers. The no-insurance scenario is associated with an average welfare loss of over 1.8% of lifetime equivalent consumption - over one and a half times that of the interventionist scenarios.

Societies do not consist of an average household, or a benevolent social planner who makes decisions on behalf of such an entity. Instead, the distribution of gains and losses from any crisis and from any policy will play a crucial part in how society decides to evaluate past actions and proposes to meet new challenges. The middle plot in Figure 15 shows the distribution of welfare losses along the pre-Covid-19 gross earnings distribution. As I have discussed throughout this section, both the impact of Covid-19 itself as well as the policy response were highly nonlinear along the dimensions of pre-pandemic incomes. Under the baseline policy welfare losses form an almost perfect downward sloping line, with an incline of -0.18% per additional decile. This suggests that the highest losses from the pandemic are concentrated amongst the highest pre-Covid-19 earners - who presumably are also best placed to bear the brunt of these losses. The uniform insurance policy is similar to the baseline policy, although it is a lot steeper at low initial income levels, which is unsurprising given the regressive nature of the transfer. Both policies strictly dominate the no-insurance option.

The last subplot plots the distribution of welfare losses by initial wealth. The baseline policy sees a steep fall in welfare as initial assets increase. Throughout this section, we saw that under the baseline calibration wealthier households appeared to be more impacted by the shock, while particularly borrowing-constrained individuals benefitted - at least relatively - from the policy response. Very wealthy households are also most likely to be affected by the consumption restrictions, leading to a large direct fall in their utility. Figure A9 in the Appendix shows the welfare effects of a counterfactual without the consumption cap. Comparing the two figures suggests that the consumption cap mainly

[34]It is no surprise that households generally do not like going through the Covid-19 pandemic, so they require negative additional consumption to remain indifferent.

affects the welfare of very rich and/or wealthy households, by forcing them to consume less than they would have liked. Under the uniform insurance policy, the distribution of welfare losses is similar, with the exception of very poor or very rich households who experience slightly lower welfare losses under the uniform policy.

Under the absence of additional benefit payments, welfare losses are exacerbated across the board, as expected but there are some interesting insights to be gleaned: Namely that there is a small proportion of households who actually preferred no intervention to any particular policy. These households are located at the very top of the wealth distribution, suggesting that these are super high net worth households who benefit from future higher interest rates brought about by a scarcity of capital.

Figure 16: Welfare effects of the Covid-19 pandemic

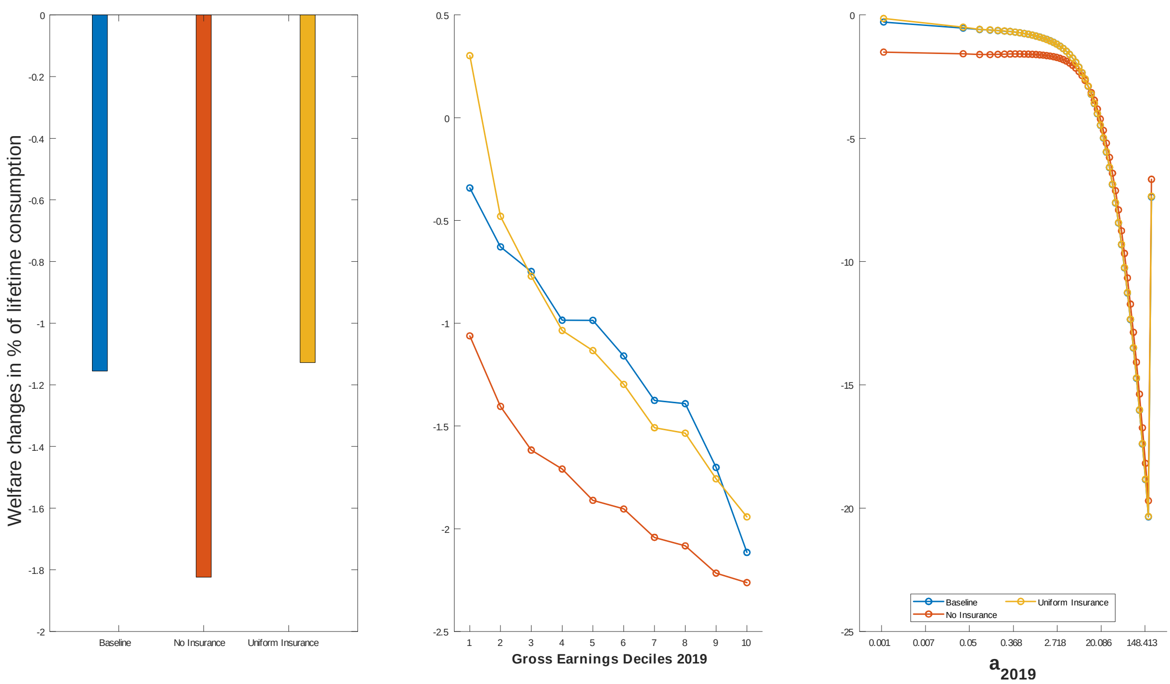

Note: Welfare effects expressed in percentages of lifetime consumption equivalents in the stationary distribution. Assets plotted on log scale.

Apart from this small plutocracy, the majority of households prefer some form of government intervention in response to the pandemic. The uniform insurance has a slight edge in terms of average welfare and is preferred in general by very poor households and also households at the very top of the pre-Covid-19 earnings distribution. In order to assess which policy choice is consistent with majority decision making I calculate the share of households that prefer each policy option. A broad majority of 73.7% of households prefer the baseline policy, while 23.2% prefer the uniform insurance option. As suggested before, only 3.1% of households prefer no additional insurance at all.

In order to provide the additional Covid-19 relief payments the UK government at the time borrowed over £300 billion, adding substantially to a large pre-existing public debt. While an assessment of the implications of different strategies to pay back the incurred debt is beyond the scope of this paper, it is evident that short of sovereign default, paying back the debt will require increased taxation at some point in the future and these costs need to be weighed against the benefits of intervention. Whilst this paper cannot speak to the optimal tax schedule it may provide some guiding thoughts on what might and what might not be advisable to consider.

In light of rising wealth inequality, governments might be tempted to introduce wealth taxation in order to pay down the national debt. Whilst such schemes can raise large revenues (e.g. Saez & Zucman (2019)) and may even increase efficiency (Guvenen et al. (2019)) it is important to ensure that these schemes are sufficiently progressive as to not disincentivise saving by households with low or moderate wealth. Throughout this paper, analysis has shown that wealth-poor households face additional challenges with respect to the accumulation of human capital, and are also much more vulnerable to unexpected economic shocks. Wealth-poor households have benefitted from the existing Covid-19 support structures precisely because they would have been ill-placed to face the impact of Covid-19 without them. This suggests that policies that support the wealth poorest households to build up some even moderate savings, would have positive consequences for skill accumulation, intergenerational mobility and resilience to economic shocks.

7 Conclusion

In this paper, I have explored the interaction between wealth and human capital in a novel heterogeneous agent incomplete market model. The model features endogenous human capital accumulation as well as savings, capturing two important dimensions of household heterogeneity. I calibrate the model to the UK economy pre-Covid-19 using microdata.

I find that there are important non-linearities in human capital investments, with workers with low levels of wealth investing considerably less in accumulating human capital than their counterparts with more wealth. This suggests the existence of low-wealth poverty traps, where individuals with low wealth struggle to increase their skillset and therefore lag behind comparable individuals with higher initial levels of wealth. Even small initial differences in household wealth can lead to sizeable and long-lasting gaps in human capital accumulation and therefore earnings inequality.

Low-wealth households are also particularly vulnerable to unexpected economic shocks, unable to draw on savings, these households reduce their investment into skills leading to persistent scarring effects in terms of human capital and incomes. The nonlinearity of the households' response to aggregate shocks means that the distribution of wealth plays an important role in determining how the economy recovers after a recession.

The model was used to analyse the economic dynamics of the distribution of human capital in the aftermath of the Covid-19 pandemic in the UK. Using a baseline calibration, the model predicts that aggregate human capital will fall significantly over the course of the next decades, leading to reduced growth prospects for the foreseeable future. Using counterfactual policy scenarios, I find that overall the Covid-19 support measures put in place throughout 2020/21, were effective in preventing a much more severe economic fallout from the pandemic and were particularly effective in protecting low-skilled and low-wealth households and workers.

The analysis thus highlights the need for decisive and targeted policy action to ensure that the UK economy can recover from the Covid-19 pandemic with a minimum of lasting damage. Measures such as targeted job support, targeted training and skills development programs and policies to promote flexible working should all be considered as potential ingredients in the UK's recipe for economic recovery. These should be complemented by measures to ensure that the gains from economic growth are widely shared and that human capital inequality does not increase further. By taking these steps, the UK can ensure that its economy is not only able to recover from the effects of the pandemic but is also better positioned to cope with future shocks and to continue to grow and develop in a sustainable way.

The model developed in the paper opens up several interesting avenues for future research: For example, it would be interesting to explore the effects of policies that aim to reduce wealth inequality and promote human capital investments, such as education policies, or Universal Basic Income. Further, much of the literature on heterogenous agent models with aggregate shocks operates under the assumption that the distribution of wealth does not matter significantly for the economy's response to aggregate shocks (see Krusell & Smith (1998)). The findings of this paper suggest that perhaps including human capital accumulation might be a fruitful avenue for making wealth inequality matter for aggregate dynamics. I leave these and other questions for future research.

BIBLIOGRAPHY

[1] Acemoglu, Daron. "Training and innovation in an imperfect labour market." The Review of Economic Studies 64, no. 3 (1997): 445-464.

[2] Acemoglu, Daron. "Obedience in the Labour Market and Social Mobility: A Socioeconomic Approach." Economica 89 (2022): S2-S37.

[3] Açıkgöz, Ömer T. "On the existence and uniqueness of stationary equilibrium in Bewley economies with production." Journal of Economic Theory 173 (2018): 18-55.

[4] Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh. "Inequality in the impact of the coronavirus shock: Evidence from real time surveys." Journal of Public economics 189 (2020): 104245.

[5] Adermon, Adrian, Mikael Lindahl, and Mårten Palme. "Dynastic human capital, inequality, and intergenerational mobility." American Economic Review 111, no. 5 (2021): 1523-48.

[6] Aghion, Philippe, and Xavier Jaravel. "Knowledge spillovers, innovation and growth." The Economic Journal 125, no. 583 (2015): 533-573.

[7] Aiyagari, S. Rao. "Uninsured idiosyncratic risk and aggregate saving." The Quarterly Journal of Economics 109, no. 3 (1994): 659-684.

[8] Almeida, Vanda, Salvador Barrios, Michael Christl, Silvia De Poli, Alberto Tumino, and Wouter van der Wielen. "The impact of COVID-19 on households income in the EU." The Journal of Economic Inequality 19, no. 3 (2021): 413-431.

[9] Angelopoulos, Angelos, Konstantinos Angelopoulos, Spyridon Lazarakis, and Apostolis Philippopoulos. "The distributional consequences of rent-seeking." Economic Inquiry 59, no. 4 (2021): 1616-1640.

[10] Angelopoulos, Konstantinos, Spyridon Lazarakis, and James Malley. "The distributional implications of asymmetric income dynamics." European Economic Review 128 (2020): 103502.

[11] Angelopoulos, Konstantinos, Spyridon Lazarakis, Rebecca Mancy, and Max Schroeder. "Pandemic-Induced Wealth and Health Inequality and Risk Exposure." (2021).

[12] Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models." Econometrica 89, no. 5 (2021): 2375-2408.

[13] Bank of England, How has Covid affected household savings? Bank of England November, 2020, Accessed 22nd February, 2023 https://www.bankofengland.co.uk/bank-overground/2020/how-has-covid-affected-household-savings (2020)

[14] Becker, Gary, and ROBERT A. Collins. "Human capital investment." (1964).

[15] Belfield, Chris, Richard Blundell, Jonathan Cribb, Andrew Hood, and Robert Joyce. "Two decades of income inequality in Britain: the role of wages, household earnings and redistribution." Economica 84, no. 334 (2017): 157-179.

[16] Ben-Porath, Yoram. "The production of human capital and the life cycle of earnings." Journal of political economy 75, no. 4, Part 1 (1967): 352-365.

[17] Bewley, Truman. "Stationary monetary equilibrium with a continuum of independently fluctuating consumers." Contributions to mathematical economics in honor of Gérard Debreu 79 (1986).

[18] Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes. "Why the apple doesn't fall far: Understanding intergenerational transmission of human capital." American economic review 95, no. 1 (2005): 437-449.

[19] Blundell, Richard, and Ben Etheridge. "Consumption, income and earnings inequality in Britain." Review of Economic Dynamics 13, no. 1 (2010): 76-102.

[20] Blundell, Richard, Monica Costa Dias, Costas Meghir, and Jonathan Shaw. "Female labor supply, human capital, and welfare reform." Econometrica 84, no. 5 (2016): 1705-1753.

[21] Blundell, Richard, Monica Costa Dias, Robert Joyce, and Xiaowei Xu. "COVID-19 and Inequalities." Fiscal studies 41, no. 2 (2020): 291-319.

[22] Blundell, Richard, Monica Costa Dias, Jonathan Cribb, Robert Joyce, Tom Waters, Thomas Wernham, and Xiaowei Xu. "Inequality and the COVID-19 crisis in the United Kingdom." Annual Review of Economics 14 (2022): 607-636.

[23] Blundell, Richard, Robert Joyce, Agnes Norris Keiller, and James P. Ziliak. "Income inequality and the labour market in Britain and the US." Journal of Public Economics 162 (2018): 48-62.

[24] Boppart, Timo, Per Krusell, and Kurt Mitman. "Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative." Journal of Economic Dynamics and Control 89 (2018): 68-92.

[25] Carroll, Christopher D. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." Economics letters 91, no. 3 (2006): 312-320.

[26] Clark, Gregory. "The son also rises." In The Son Also Rises. Princeton University Press, 2014.

[27] Crafts, Nicholas, and Terence C. Mills. "Is the UK productivity slowdown unprecedented?." National Institute Economic Review 251 (2020): R47-R53.

[28] Davenport, A., Joyce, R., Rasul, I., and T. Waters. Spending and saving during the COVID-19 crisis: evidence from bank account data, IFS: Briefing note BN308, The Institute for Fiscal Studies. (2020)

[29] Dinerstein, Michael, Rigissa Megalokonomou, and Constantine Yannelis. Human capital depreciation. No. w27925. National Bureau of Economic Research, 2020.

[30] Dossche, Maarten, and Stylianos Zlatanos. "COVID-19 and the increase in household savings: precautionary or forced?." Economic Bulletin Boxes 6 (2020).

[31] Dräger, Jascha. "The role of parental wealth in children's educational pathways in Germany." European Sociological Review 38, no. 1 (2022): 18-36.

[32] Faccini, Renato, Stephen Millard, and Francesco Zanetti. "Wage rigidities in an estimated dynamic, stochastic, general equilibrium model of the UK labour market." The Manchester School 81 (2013): 66-99.

[33] Fan, Xiaodong, Ananth Seshadri, and Christopher R. Taber. Estimation of a life-cycle model with human capital, labor supply and retirement. No. w29905. National Bureau of Economic Research, 2022.

[34] Gourieroux, Christian, Monfort Gourieroux, Alain Monfort, and Director Alain Monfort. Simulation-based econometric methods. Oxford university press, 1996.

[35] Guvenen, Fatih, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen. Use it or lose it: Efficiency gains from wealth taxation. No. w26284. National Bureau of Economic Research, 2019.

[36] Guvenen, Fatih, Burhanettin Kuruscu, and Serdar Ozkan. "Taxation of human capital and wage inequality: A cross-country analysis." Review of Economic Studies 81, no. 2 (2014): 818-850.

[37] Guvenen, Fatih, Serdar Ozkan, and Jae Song. "The nature of countercyclical income risk." Journal of Political Economy 122, no. 3 (2014): 621-660.

[38] Hacıoğlu-Hoke, Sinem, Diego R. Känzig, and Paolo Surico. "The distributional impact of the pandemic." European Economic Review 134 (2021): 103680.

[39] Harari, Daniel, Matthew Keep, and Philip Brien. "Coronavirus: economic impact." Briefing paper 8866 (2021).

[40] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. "Optimal tax progressivity: An analytical framework." The Quarterly Journal of Economics 132, no. 4 (2017): 1693-1754.

[41] Heckman, James J., Lance Lochner, and Christopher Taber. "Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents." Review of economic dynamics 1, no. 1 (1998): 1-58.

[42] Heer, Burkhard, and Alfred Maussner. Dynamic general equilibrium modeling: computational methods and applications. Springer Science & Business Media, 2009.

[43] Hubmer, Joachim, Per Krusell, and Anthony A. Smith Jr. "Sources of US wealth inequality: Past, present, and future." NBER Macroeconomics Annual 35, no. 1 (2021): 391-455.

[44] Huckfeldt, C., 2022. Understanding the scarring effect of recessions. American Economic Review, 112(4), pp.1273-1310.

[45] Huggett, Mark. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." Journal of economic Dynamics and Control 17, no. 5-6 (1993): 953-969.

[46] Huggett, Mark, Gustavo Ventura, and Amir Yaron. "Sources of lifetime inequality." American Economic Review 101, no. 7 (2011): 2923-2954.

[47] Imrohoroğlu, Ayşe. "Cost of business cycles with indivisibilities and liquidity constraints." Journal of Political economy 97, no. 6 (1989): 1364-1383.

[48] Kaplan, Greg, and Giovanni L. Violante. "A model of the consumption response to fiscal stimulus payments." Econometrica 82, no. 4 (2014): 1199-1239.

[49] Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. "Monetary policy according to HANK." American Economic Review 108, no. 3 (2018): 697-743.

[50] Karagiannaki, Eleni. "The effect of parental wealth on children's outcomes in early adulthood." The Journal of Economic Inequality 15 (2017): 217-243.

[51] Kuhn, Moritz. Welfare analysis with permanent income shocks. Working paper, 2008, University of Mannheim, 2008.

[52] Krebs, Tom. "Human capital risk and economic growth." The Quarterly Journal of Economics 118, no. 2 (2003): 709-744.

[53] Krusell, Per, and Anthony A. Smith, Jr. "Income and wealth heterogeneity in the macroeconomy." Journal of political Economy 106, no. 5 (1998): 867-896.

[54] Lachance, Marie-Eve. "Optimal onset and exhaustion of retirement savings in a life-cycle model." Journal of Pension Economics & Finance 11, no. 1 (2012): 21-52.

[55] Lazear, Edward P. Human wealth and human capital. No. w0097. National Bureau of Economic Research, 1975.

[56] Lochner, Lance, and Alexander Monge-Naranjo. "Credit constraints in education." Annu. Rev. Econ. 4, no. 1 (2012): 225-256.

[57] Lovenheim, Michael F. "The effect of liquid housing wealth on college enrollment." Journal of Labor Economics 29, no. 4 (2011): 741-771.

[58] Lundborg, Petter, Martin Nordin, and Dan Olof Rooth. "The intergenerational transmission of human capital: the role of skills and health." Journal of Population Economics 31, no. 4 (2018): 1035-1065.

[59] Ma, Xiao, Alejandro Nakab, and Daniela Vidart. Human capital investment and development: The role of on-the-job training. No. 2021-10. University of Connecticut, Department of Economics, 2021.

[60] Marmot, Michael, and Jessica Allen. "COVID-19: exposing and amplifying inequalities." J Epidemiol Community Health 74, no. 9 (2020): 681-682.

[61] Marmot, Michael, Jessica Allen, Peter Goldblatt, Eleanor Herd, and Joana Morrison. "Build Back Fairer." London: Institute of Health Equity, Available online: https://www. health. org. uk/sites/default/files/upload/publications/2020/Build-back-fairer-the-COVID-19-Marmot-review. pdf (accessed on 16 March 2022) (2021).

[62] Miescu, Mirela, and Raffaele Rossi. "COVID-19-induced shocks and uncertainty." European Economic Review 139 (2021): 103893.

[63] Mincer, Jacob A. "The human capital earnings function." In Schooling, experience, and earnings, pp. 83-96. NBER, 1974.

[64] Ouyang, Min. "The scarring effect of recessions." Journal of Monetary Economics 56, no. 2 (2009): 184-199.

[65] Ríos-Rull, José-Víctor. "Life-cycle economies and aggregate fluctuations." The Review of Economic Studies 63, no. 3 (1996): 465-489.

[66] Saez, Emmanuel, and Gabriel Zucman. "Progressive wealth taxation." Brookings Papers on Economic Activity 2019, no. 2 (2019): 437-533.

[67] Solon, Gary. "What do we know so far about multigenerational mobility?." The Economic Journal 128, no. 612 (2018): F340-F352.

[68] Stantcheva, Stefanie. "Inequalities in the times of a pandemic." Economic Policy 37, no. 109 (2022): 5-41.

[69] Telemo, Paul. "Essays on wage determinants in the long and the short run." (2022).

[70] Tenreyro. Response to the Covid-19 pandemic: UK and US experiences, Bank of England speech. (2021)

[71] Quadrini, Vincenzo, and José-Víctor Ríos-Rull. "Inequality in macroeconomics." In Handbook of income distribution, vol. 2, pp. 1229-1302. Elsevier, 2015.

[72] Whitley, Elise, Ian J. Deary, Stuart J. Ritchie, G. David Batty, Meena Kumari, and Michaela Benzeval. "Variations in cognitive abilities across the life course: Cross-sectional evidence from Understanding Society: The UK Household Longitudinal Study." Intelligence 59 (2016): 39-50.

[73] Young, Eric R. "Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations." Journal of Economic Dynamics and Control 34, no. 1 (2010): 36-41.

8 Appendix A

Figure A1: Impact of alternative policy interventions on gross and net earnings

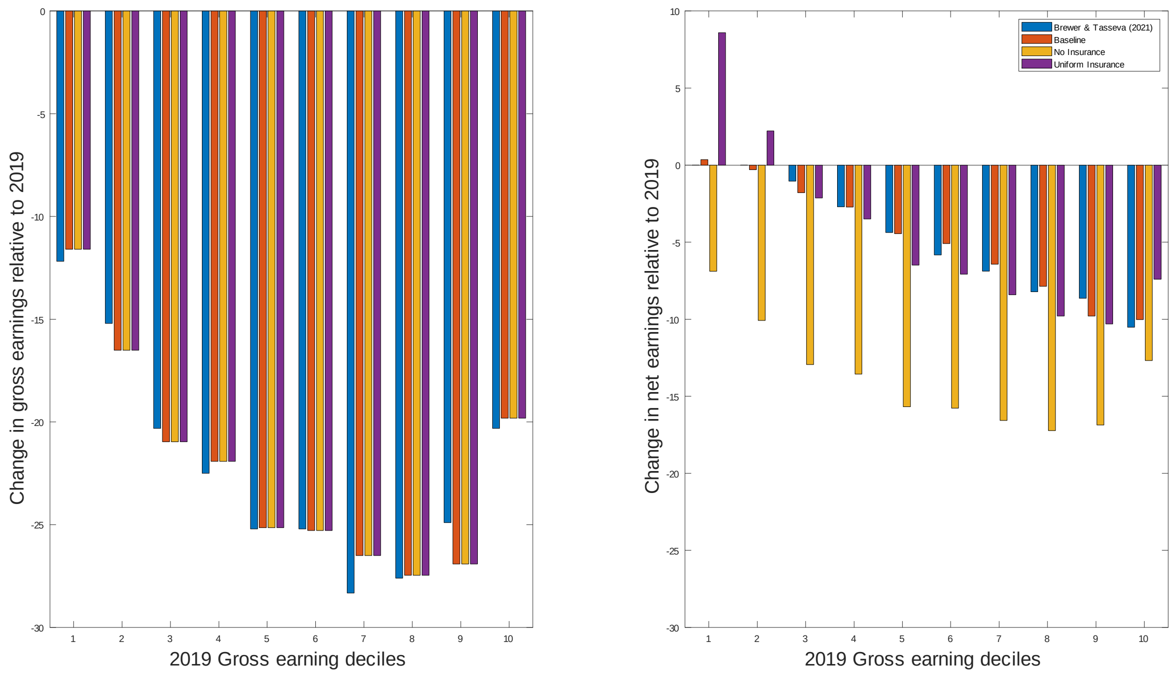

Note: Data points transcribed from Brewer & Tasseva (2021), Figure 1.

Figure A2: Response of aggregate variables to Covid-19

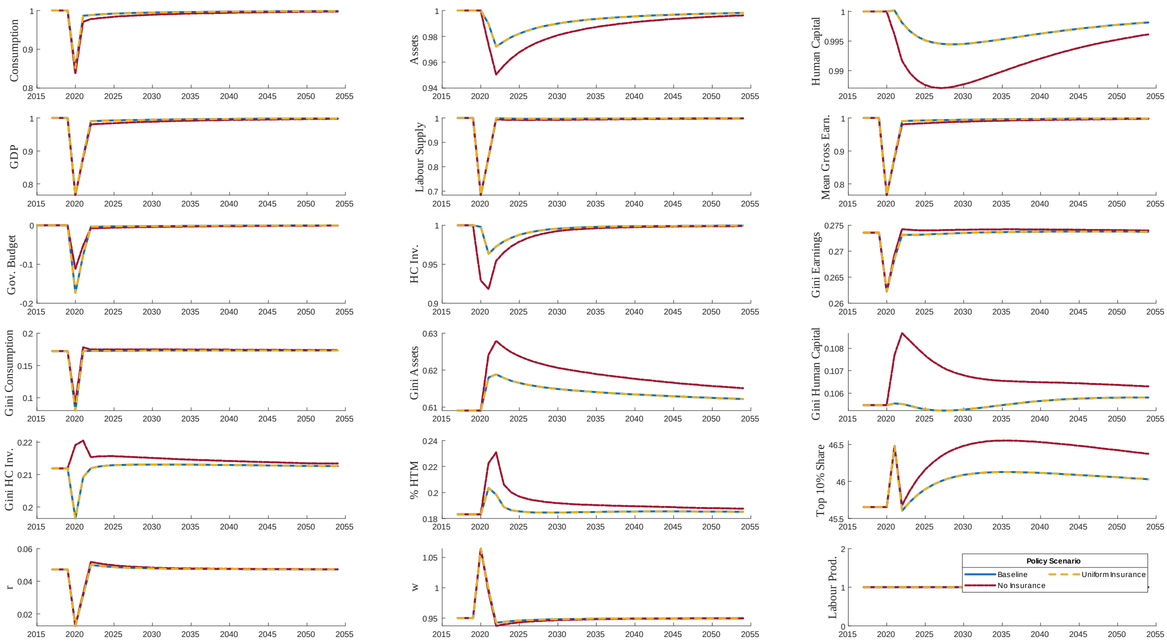

Note: Baseline, no insurance and uniform insurance Covid-19 shock calibration. Series normalised to expected progression in stationary equilibrium.

Figure A3: Consumption response to Covid-19, by wealth quantiles

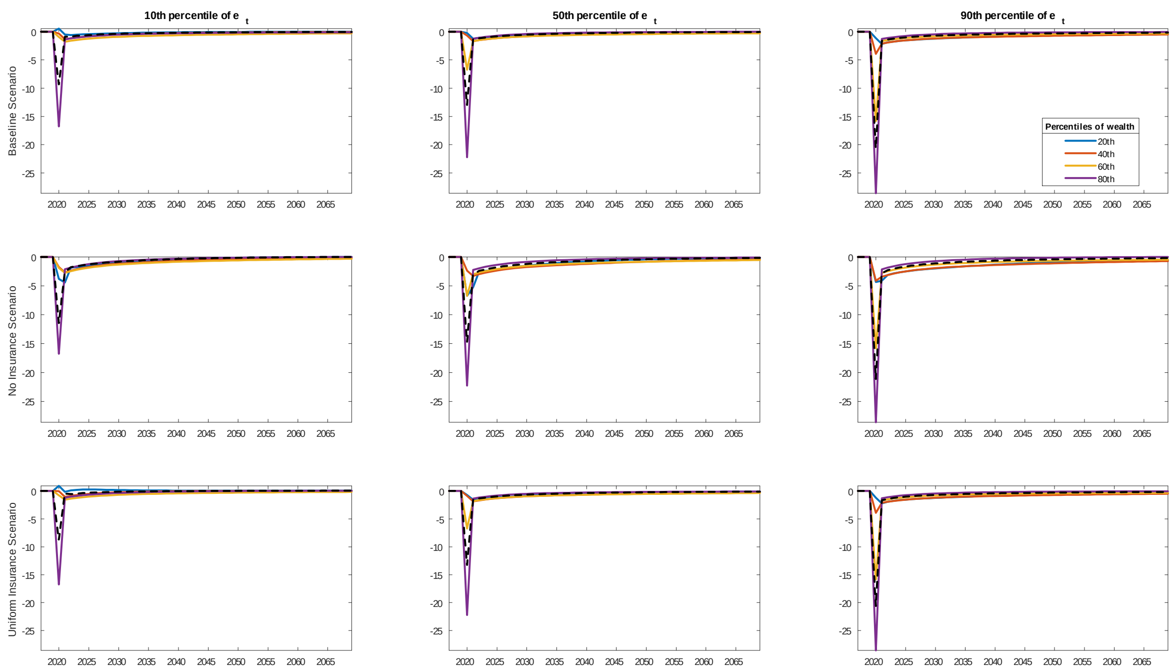

Note: Series normalised to expected progression in stationary equilibrium. Workers begin with human capital equal to the mean level of human capital within their wealth-productivity percentile. See Table 8 for details.

Figure A4: Human capital investment response to Covid-19, by wealth quantiles

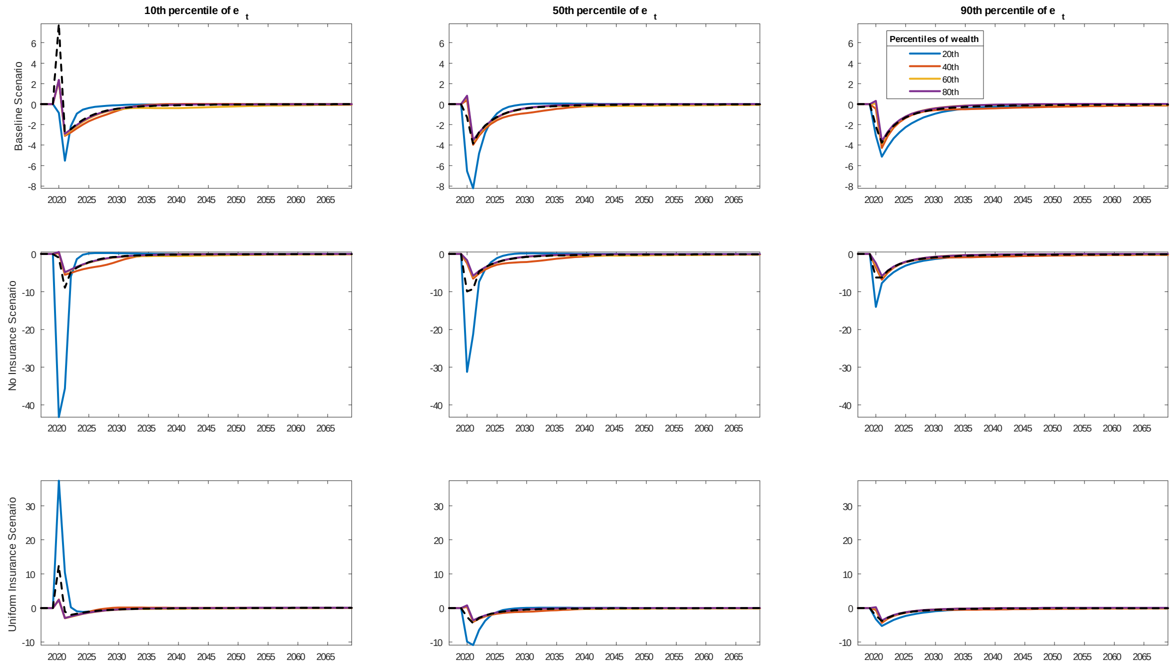

Note: Series normalised to expected progression in stationary equilibrium. Workers begin with human capital equal to the mean level of human capital within their wealth-productivity percentile. See Table 8 for details.

Figure A5: Comparison of positive and negative TFP shocks

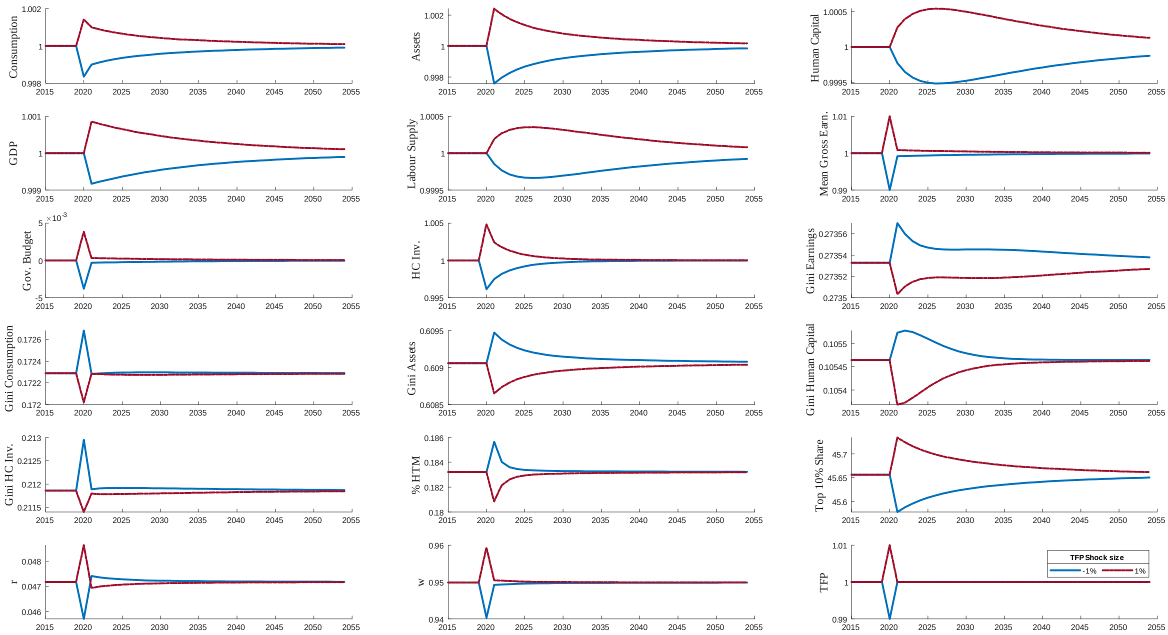

Note: Response of aggregate variables to different TFP shocks.

Figure A6: Response of aggregate variables to Covid-19

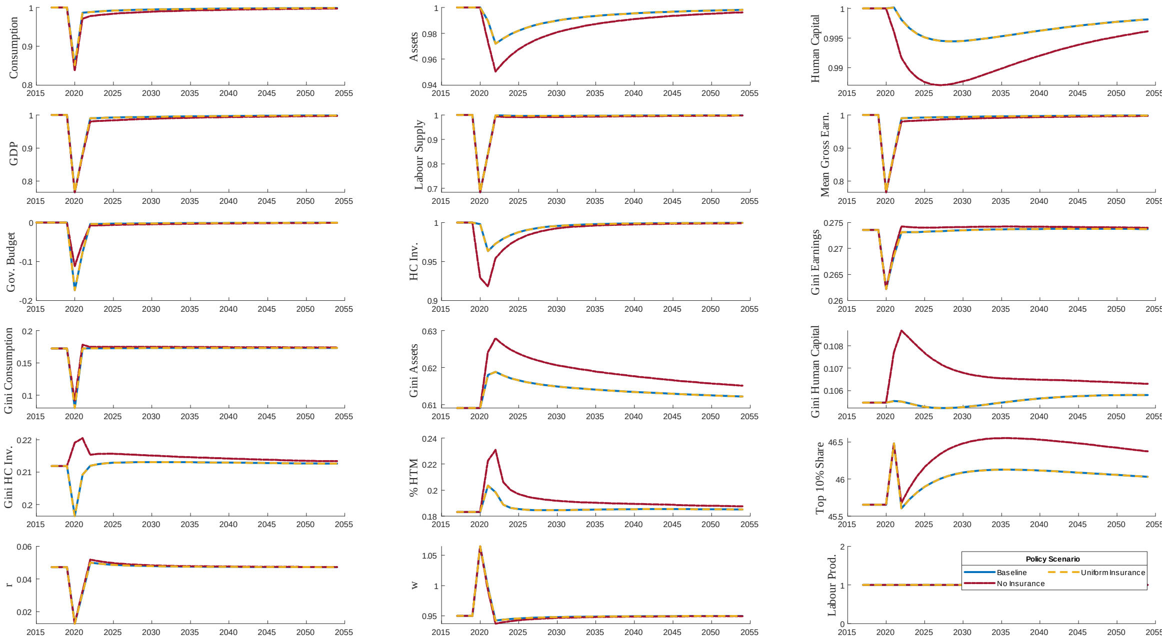

Note: Baseline, no insurance and uniform insurance Covid19 shock calibration. Series normalised to expected progression in stationary equilibrium.

Figure A7: Impact of consumption cap on household expenditure and savings

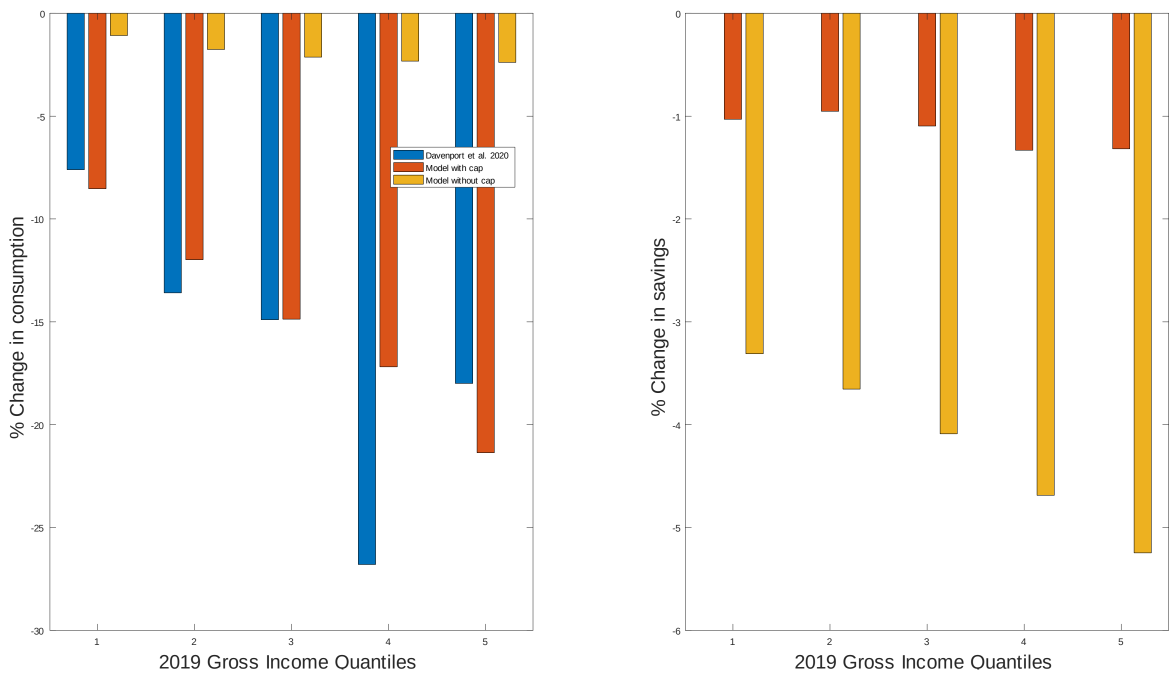

Note: Data points transcribed from Davenport et al. (2020), Figure 4.1.

Baseline model with and without consumption restrictions.

Changes relative to averages in pre-Covid-19 period.

Figure A8: Baseline model with and without consumption restrictions

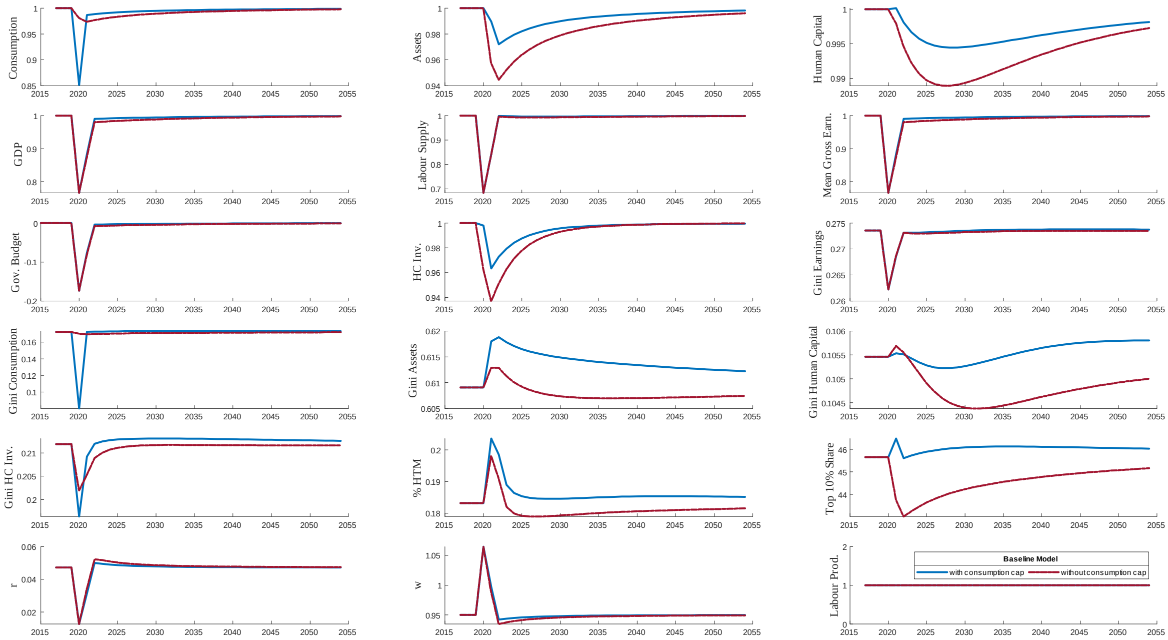

Note: Baseline Covid-19 shock calibration.

Series normalised to expected progression in stationary equilibrium.

Figure A9: Welfare effects of the Covid-19 pandemic, no consumption cap

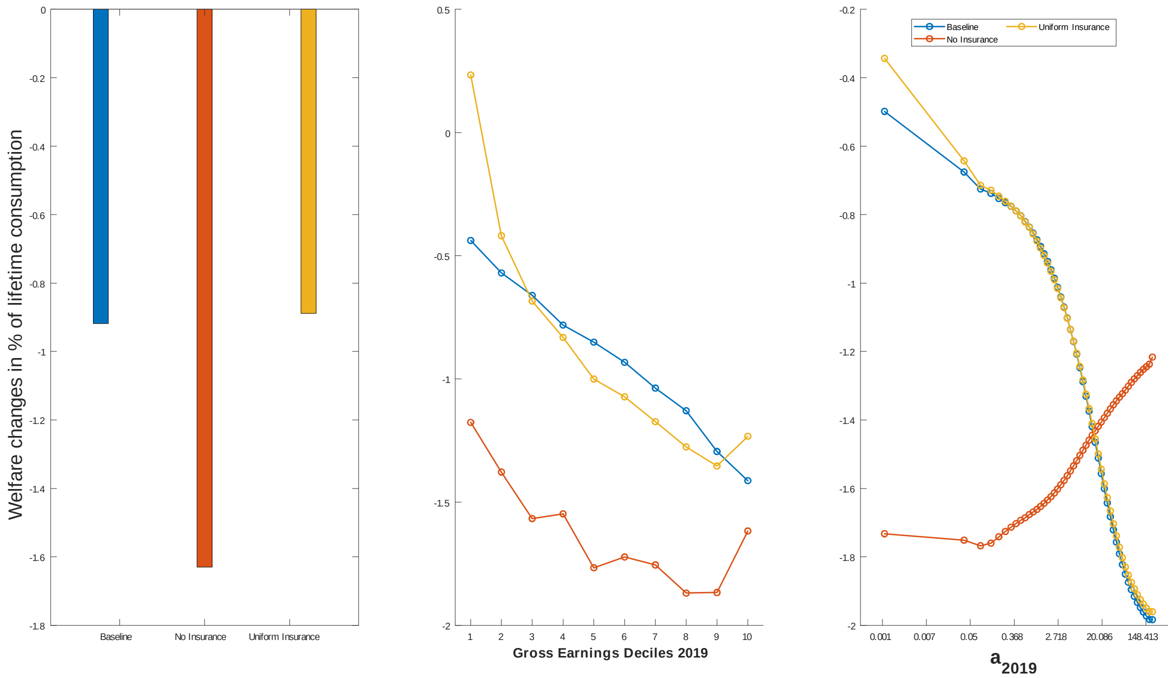

Note: Welfare effects expressed in percentages of lifetime consumption equivalents in the stationary distribut
Assets plotted on log scale.

9 Appendix B

9.1 EGM with 2 state variables

This algorithm is an extended version of Caroll's EGM algorithm (Caroll (2006)). It's based on the extension of the method presented in Auclert et al. (2021), but I have adapted it for the case of human capital.

The problem of the household

The households problem is given by:

$$V\left(a_t,h_t,e_t\right)=\max_{\{c_t,a_{t+1},x_t\}}\left\{\frac{\left(c_t\right)^{1-\sigma}}{1-\sigma}+\beta E\left[V\left(a_{t+1},h_{t+1},e_{t+1}\right)\mid e_t\right]\right\}$$

$$s.t.$$

$$c_t+a_{t+1}+x_t=(1+r_t)a_t+g(h_t,e_t,w_t,\tau_1,\tau_2)$$

$$h_{t+1}=h_t\delta_t+\chi x_t^{\nu}$$

$$g(h_t,e_t,w_t,\tau_1,\tau_2)=\tau_1(w_te_t\log(1+h_t))^{1-\tau_2}$$

$$c_t\geq 0, h_t\geq 0,\ a_t\geq 0, x_t\geq 0$$

$$\text{endogenous states :}h_t,a_t$$

$$\text{exogenous states :}\delta_t,e_t$$

$$\text{choice variables :}c_t,a_{t+1},h_{t+1},x_t$$

First, rewrite the human capital investment function in implicit form:

$$x_t=m(h_{t+1},h_t)=\left[\frac{(h_{t+1}-\delta_th_t)}{\chi}\right]^{\frac{1}{\upsilon}}.$$

To impose the non-negativity constraint on $x_t\geq 0$, further impose:

$$x_t=m(h_{t+1},h_t)=\left[\frac{abs\left(h_{t+1}-\delta_th_t\right)}{\chi}\right]^{\frac{1}{\upsilon}}.$$

Then substitute out consumption and human capital investment to rewrite the bellman equation:

$$V\left(a_t,h_t,e_t\right)\max_{\{h_{t+1},a_{t+1}\}}\left\{\begin{array}{c}u(g(h_t,e_t,w_t,\tau_1,\tau_2)+(1+r)a_t-a_{t+1}-m(h_{t+1},h_t))\\+\lambda_t(a_{t+1}-\underline{a})+\mu_t(h_{t+1}-\delta_th_t)+\beta E\left[V\left(a_{t+1},h_{t+1},e_{t+1}\right)\mid e_t\right]\end{array}\right\}$$

The first order conditions are:

$$u'(c_t)=\lambda_t+\beta E\left[\partial_{a_{t+1}}V\left(a_{t+1},h_{t+1},e_{t+1}\right)\mid e_t\right]$$

and

$$u'(c_t)m_1(h_{t+1},h_t)=\mu_t+\beta E\left[\partial_{h_{t+1}}V\left(a_{t+1},h_{t+1},e_{t+1}\right)\mid e_t\right]$$

and the envelope conditions are:

$$\partial_{a_t} V\left(a_t, h_t, e_t\right)=(1+r) u^{\prime}\left(c_t\right)$$

and

$$\partial_{h_t} V\left(a_t, h_t, e_t\right)=\left(g_1\left(h_t, e_t, w_t, \tau_1, \tau_2\right)-m_2\left(h_{t+1}, h_t\right)\right) u^{\prime}\left(c_t\right)$$

Algorithm

The algorithm follows Auclert et al. (2021).

1. Initial guess: Make two guesses for the derivatives of the value function $V_a(a', h', e')$ and $V_h(a', h', e')$.

2. Common $e' \longrightarrow e$. Calculate the discounted expectation using the exogenous transition matrix Π:

$$\begin{aligned} W_a(a', h', e) &= \beta \Pi V_a(a', h', e') \\ W_h(a', h', e) &= \beta \Pi V_h(a', h', e') \end{aligned}$$

3. Unconstrained $h' \longrightarrow h$. Assume neither constraint is binding, then $\lambda_t = \mu_t = 0$, and the FOCs are

$$\begin{aligned} c^{-\sigma} &= W_a(a', h', e) \\ c^{-\sigma} m_1\left(h', h\right) &= W_h(a', h', e) \end{aligned}$$

   Divide each side to get:

$$\frac{1}{m_1\left(h', h\right)} = \frac{W_a(a', h', e)}{W_h(a', h', e)}.$$

   Define $F(h, h', a', e)$ by setting the previous expression equal to 0:

$$F(h, h', a', e) = \frac{W_a(a', h', e)}{W_h(a', h', e)} - \frac{1}{m_1\left(h', h, e\right)} = 0.$$

   Note that $m_1\left(h', h\right) > 0$ **iff** $\mu_t = 0$. $F(h, h', a', e)$ is a mapping that characterizes $h'(a', h, e)$. Use this mapping to map $W_a(a', h', e)$ into $W_a(a', h, e)$ using linear interpolation. Compute consumption as

$$c(a', h, e) = W_a(a', h, e)^{-\frac{1}{\sigma}}.$$

4. Unconstrained $a' \longrightarrow a$. Now use the budget constraint and $h'(a', h, e)$ and $c(a', h, e)$

from the previous step to obtain:

$$a(a', h, e) = \frac{a' + c(a', h, e) + m\left(h'(a', h, e), h\right) - g(h, e, w)}{(1 + r)}.$$

Invert this function via interpolation to get $a'(a, h, e)$ and the same interpolation weights can be used to do $h'(a', h, e) \longrightarrow h'(a, h, e)$.

5. Liquidity constrained $h' \longrightarrow h$. This branch proceeds analogously to the unconstrained case. Assuming that the liquidity constraint is binding, $\lambda_t > 0$, but not the minimum human capital investment constraint, $\mu_t = 0$. The FOCs become:

$$\begin{aligned} c^{-\sigma} &= \lambda + W_a(\underline{a}, h', e) \\ c^{-\sigma} m_1\left(h', h\right) &= W_h(\underline{a}, h', e). \end{aligned}$$

For scaling purposes, define $\kappa \equiv \lambda / W_a(\underline{a}, h', e)$ and rewrite the first FOC as:

$$c^{-\sigma} = (1 + \kappa) W_a(\underline{a}, h', e)$$

Divide side by side to get:

$$\frac{1}{m_1\left(h', h, e\right)} = \frac{(1 + \kappa) W_a(\underline{a}, h', e)}{W_h(\underline{a}, h', e)}$$

$$F(h, h', \kappa, e) = \frac{(1 + \kappa) W_a(\underline{a}, h', e)}{W_h(\underline{a}, h', e)} - \frac{1}{m_1\left(h', h, e\right)} = 0.$$

Solve this for $h'(\kappa, h, e)$ and compute consumption as:

$$c(\kappa, h, e) = \left[(1 + \kappa) W_a(a', h, e)\right]^{-\frac{1}{\sigma}}.$$

6. Liquidity constrained $\kappa \longrightarrow a$. Now using $h'(\kappa, h, e)$ and $c(\kappa, h, e)$ from the previous step, use the budget constraint to obtain

$$a(\kappa, h, e) = \frac{\underline{a} + c(\kappa, h, e) + m\left(h'(\kappa, h, e), h\right) - g(h, e, w)}{(1 + r)}$$

Invert this function via interpolation to get $\kappa(a, h, e)$. The same interpolation weights can be used to map $h'(\kappa, h, e)$ into $h'(a, h, e)$. We already know that $a'(a, h, e) = \underline{a}$.

7. Binding human capital investment constraint. Suppose that both constraints bind

$\lambda_t \neq 0, \mu_t \neq 0$. Then the FOCs are:

$$c^{-\sigma} = \lambda + W_a(\underline{a}, h', e)$$
$$c^{-\sigma} m_1\left(h', h\right) = \mu + W_h(\underline{a}, h', e).$$

We already know that $m_1\left(h', h\right) > 0$ **iff** $\mu_t = 0$. We also know that $m_1\left(h', h\right) \geq 0$.It follows that if $\mu_t > 0$ then $m_1\left(h', h\right) = 0$. So from the second FOC:

$$\mu = -W_h(\underline{a}, h', e).$$

This means that the gradient of the value function with respect to human capital is negative, which is inconsistent with the economic logic of the problem. Hence, the constraint of human capital investment never binds, except for the trivial case when the household has zero resources. Intuitively, convex human capital investment function implies that it is always profitable for the agent to invest an incremental amount into human capital.

8. Update guesses. The final $a'(a, h, e)$ is the element-wise maximum of its unconstrained and liquidity-constrained counterparts. Replace the unconstrained $h'(a, h, e)$ with constrained one at the exact same points. Compute consumption from the budget constraint as

$$c(a, h, e) = (1 + r)a + g(h, e, w) - m\left(h'(a, h, e), h\right) - a'(a, h, e)$$

Finally use the envelope conditions and to update the guesses

$$V_a(a, h, e) = (1 + r)c(a, h, e)^{-\sigma}$$

$$V_h(a, h, e) = \left[g_1(h, e, w) - m_2\left(h'(a, h, e), h\right)\right] c(a, h, e)^{-\sigma}$$

Go back to step 2, repeat until convergence.

9.2 Algorithm with consumption cap

Assume that occasionally - due to external events - agents are restricted in their consumption behaviour. This is implemented by invoking a consumption cap, and a penalty function that penalizes households for consuming above the capped level. The penalty

function is given by:
$$\frac{\xi}{3}\{\min(\bar{c}-c_t),0\}^3$$

where ξ is a parameter that determines the strength of the penalty. Households can still consume beyond the threshold, but they do incur additional utility costs when they do. In times when the consumption cap is not active $\bar{c}$ is just set to $+\infty$.

We skip straight to the bellman, where we insert the penalty function:

$$V\left(a_t,h_t,e_t\right)\max_{\{h_{t+1},a_{t+1}\}}\left\{\begin{array}{c} u(c_t)+\frac{\xi}{3}\{\min(\bar{c}-c_t),0\}^3 \\ +\lambda_t(a_{t+1}-\underline{a})+\mu_t(h_{t+1}-\delta_t h_t)+\beta E\left[V\left(a_{t+1},h_{t+1},e_{t+1}\right)\mid e_t\right] \end{array}\right\}$$

The first order conditions are now:

$$(u'(c_t)-\min\{\bar{c}-c,0\}^2)=\lambda_t+\beta E\left[\partial_{a_{t+1}}V\left(a_{t+1},h_{t+1},e_{t+1}\right)\mid e_t\right]$$

and

$$(u'(c_t)-\min\{\bar{c}-c,0\}^2)m_1\left(h',h\right)=\mu_t+\beta E\left[\partial_{h_{t+1}}V\left(a_{t+1},h_{t+1},e_{t+1}\right)\mid e_t\right]$$

and the envelope conditions are:

$$\partial_{a_t}V\left(a_t,h_t,e_t\right)=(1+r)(u'(c_t)-\min\{\bar{c}-c,0\}^2)$$

and

$$\partial_{h_t}V\left(a_t,h_t,e_t\right)=\left(g_1(h_t,e_t,w_t,\tau_1,\tau_2)-m_2(h_{t+1},h_t)\right)\left(u'(c_t)-\min\{\bar{c}-c,0\}^2\right)$$

Algorithm with consumption cap

The algorithm follows Auclert et al. (2021).

1. Initial guess: Make two guesses for the derivatives of the value function $V_a(a',h',e')$ and $V_h(a',h',e')$.

2. Common $e'\longrightarrow e$. Calculate the discounted expectation using the exogenous transition matrix Π:

$$W_a(a',h',e)=\beta\Pi V_a(a',h',e')$$
$$W_h(a',h',e)=\beta\Pi V_h(a',h',e')$$

3. Unconstrained $h' \longrightarrow h$. Assume neither constraint is binding, then $\lambda_t = \mu_t = 0$, and the FOCs are

$$(c^{-\sigma} - \min\{\bar{c} - c, 0\}^2) = W_a(a', h', e)$$
$$(c^{-\sigma} - \min\{\bar{c} - c, 0\}^2)m_1(h', h) = W_h(a', h', e)$$

Divide each side to get:

$$\frac{1}{m_1(h', h)} = \frac{W_a(a', h', e)}{W_h(a', h', e)}.$$

Define $F(h, h', a', e)$ by setting the previous expression equal to 0:

$$F(h, h', a', e) = \frac{W_a(a', h', e)}{W_h(a', h', e)} - \frac{1}{m_1(h', h, e)} = 0.$$

Note that $m_1(h', h) > 0$ **iff** $\mu_t = 0$. $F(h, h', a', e)$ is a mapping that characterizes $h'(a', h, e)$. Use this mapping to map $W_a(a', h', e)$ into $W_a(a', h, e)$ using linear interpolation. Compute consumption by solving the nonlinear equation:

$$c(a', h, e) = (W_a(a', h, e) + \min\{\bar{c} - c, 0\}^2)^{-\frac{1}{\sigma}}.$$

4. Unconstrained $a' \longrightarrow a$. Now use the budget constraint and $h'(a', h, e)$ and $c(a', h, e)$ from the previous step to obtain:

$$a(a', h, e) = \frac{a' + c(a', h, e) + m(h'(a', h, e), h) - g(h, e, w)}{(1 + r)}.$$

Invert this function via interpolation to get $a'(a, h, e)$ and the same interpolation weights can be used to do $h'(a', h, e) \longrightarrow h'(a, h, e)$.

5. Liquidity constrained $h' \longrightarrow h$. This branch proceeds analogously to the unconstrained case. Assuming that the liquidity constraint is binding, $\lambda_t > 0$, but not the minimum human capital investment constraint, $\mu_t = 0$. The FOCs become:

$$(c^{-\sigma} - \min\{\bar{c} - c, 0\}^2) = \lambda + W_a(\underline{a}, h', e)$$
$$(c^{-\sigma} - \min\{\bar{c} - c, 0\}^2)m_1(h', h) = W_h(\underline{a}, h', e).$$

For scaling purposes, define $\kappa \equiv \lambda / W_a(\underline{a}, h', e)$ and rewrite the first FOC as:

$$(c^{-\sigma} - \min\{\bar{c} - c, 0\}^2) = (1 + \kappa)W_a(\underline{a}, h', e)$$

Divide side by side to get:

$$\frac{1}{m_1\left(h',h,e\right)} = \frac{(1+\kappa)W_a(\underline{a},h',e)}{W_h(\underline{a},h',e)}$$

$$F(h,h',\kappa,e) = \frac{(1+\kappa)W_a(\underline{a},h',e)}{W_h(\underline{a},h',e)} - \frac{1}{m_1\left(h',h,e\right)} = 0.$$

Solve this for $h'(\kappa,h,e)$ and compute consumption by solving:

$$c(\kappa,h,e) = \left[(1+\kappa)W_a(a',h,e) + \min\{\bar{c}-c,0\}^2\right]^{-\frac{1}{\sigma}}.$$

6. Liquidity constrained $\kappa \longrightarrow a$. Now using $h'(\kappa,h,e)$ and $c(\kappa,h,e)$ from the previous step, use the budget constraint to obtain

$$a(\kappa,h,e) = \frac{\underline{a} + c(\kappa,h,e) + m\left(h'(\kappa,h,e),h\right) - g(h,e,w)}{(1+r)}$$

   Invert this function via interpolation to get $\kappa(a,h,e)$. The same interpolation weights can be used to map $h'(\kappa,h,e)$ into $h'(a,h,e)$. We already know that $a'(a,h,e) = \underline{a}$.

7. Binding human capital investment constraint. Suppose that both constraints bind $\lambda_t \neq 0, \mu_t \neq 0$. Then the FOCs are:

$$c^{-\sigma} = \lambda + W_a(\underline{a},h',e)$$
$$c^{-\sigma}m_1\left(h',h\right) = \mu + W_h(\underline{a},h',e).$$

   We already know that $m_1\left(h',h\right) > 0$ **iff** $\mu_t = 0$. We also know that $m_1\left(h',h\right) \geq 0$.It follows that if $\mu_t > 0$ then $m_1\left(h',h\right) = 0$. So from the second FOC:

$$\mu = -W_h(\underline{a},h',e).$$

   This means that the gradient of the value function with respect to human capital is negative, which is inconsistent with the economic logic of the problem. Hence, the constraint of human capital investment never binds, except for the trivial case when the household has zero resources. Intuitively, convex human capital investment function implies that it is always profitable for the agent to invest an incremental amount into human capital.

8. Update guesses. The final $a'(a,h,e)$ is the element-wise maximum of its unconstrained and liquidity-constrained counterparts. Replace the unconstrained $h'(a,h,e)$ with constrained one at the exact same points. Compute consumption from the budget constraint as

$$c(a,h,e) = (1+r)a + g(h,e,w) - m\left(h'(a,h,e),h\right) - a'(a,h,e)$$

Finally use the envelope conditions and to update the guesses

$$V_a(a,h,e) = (1+r)(c(a,h,e)^{-\sigma} - \min\{\bar{c}-c,0\}^2)$$

$$V_h(a,h,e) = \left[g_1(h,e,w) - m_2\left(h'(a,h,e),h\right)\right] c(a,h,e)^{-\sigma}(c(a,h,e)^{-\sigma} - \min\{\bar{c}-c,0\}^2)$$

Go back to step 2, repeat until convergence.

9.3 Dynamic transitions

To obtain the dynamic transition paths following the shock, I follow the MIT-shock method (c.f. Boppart et al. (2017)):

1. Choose a time T at which the economy has presumably reached a steady state.[35]
2. Solve for the stationary equilibrium at T.
3. Guess a transition path for all aggregate variables, and obtain the relevant prices given the aggregates.
4. Given the guess for the transition path, solve the policy functions backwards from $t = T - 1$.
5. Calculate the transition matrix for the joint state in every time period, using the policy functions obtained in step **4.** Iterate the joint distribution forward, starting with the initial stationary distribution.
6. Calculate the implied evolution of the aggregate states at each point in time.
7. Compare the path of the aggregates with the initial guess and update the guess until convergence is reached.

For the simulation of the joint distribution I use a histogram approach with linear interpolation between gridpoints (c.f. Angelopoulos et al. (2021)).

[35] For my Covid-19 application I set T= 500. The MIT shock dynamics use T=150, since the shock is relatively small.

9.4 Calibration of the Covid-19 shock

I calibrate the impact of Covid-19 on exogenous human capital productivity and the corresponding additional benefit payments as follows:

1. I obtain the percentage reductions in gross earnings and post all taxes and transfers net labour incomes for the 10 deciles of the pre-Covid-19 gross earnings distribution from Figure 1 in Brewer & Tasseva (2021).

2. I assign every state combination of (a_t, h_t, e_t) to one gross earnings decile using the stationary distribution of households across states and the stationary wage rate.

3. I initiate the vector κ to a random guess.

4. I parameterize both e_t^{shock} and b_t^{shock} as polynomial functions of e_t :

$$e_t^{shock}(e_t) = \kappa_0 + \kappa_1 e_t^{0.5} + \kappa_2 e_t + \kappa_3 e_t^2 \tag{2.25}$$

$$b_t^{shock}(e_t) = \max(0, \kappa_4 + \kappa_5 e_t^{0.5} + \kappa_6 e_t + \kappa_7 e_t^2) \tag{2.26}$$

5. Given the current value of $e_t^{shock}(e_t)$ I calculate the aggregate labour supply, using the fact that the distribution of human capital is fixed at the arrival of the shock. Using the aggregate supply of capital which is also fixed at the beginning of the shock period, I calculate the wage rate $w_t = (1-\alpha)A_t\left(\frac{K_t}{L_t}\right)^\alpha$.

6. Using w_t and $e_t^{shock}(e_t)$ I calculate the implied average gross income for each initial gross earnings decile. Then I use the current value of $b_t^{shock}(e_t)$ to calculate the corresponding net labour income values.

7. Evaluate the fit of the approximation, by comparing the sum of squared percentage deviations of the gross and net earnings losses implied by the model relative to the numbers obtained in step 1.

8. Return to step 3, update κ using a nonlinear solver, and repeat until the desired level of convergence is achieved.

10 Appendix C

10.1 Data

Understanding Society (UnSoc, ISER (2020)) is a comprehensive longitudinal survey of about 40,000 households in the UK. It investigates a broad spectrum of social, economic and behavioural aspects, making it pertinent to a variety of researchers and decision-makers. Data gathering for each wave spans 24 months, and the first wave ran from January 2009 to January 2011. Although the waves overlap, the same respondents are interviewed at approximately the same time every year; no one is interviewed twice within a wave or a calendar year (see Knies (2018)). The UnSoc data in this paper refer to the free "End User Licence" versions of the datasets, SN-6614. In this paper I use waves 1-9, which means that the last wave ends just before the Covid-19 pandemic.

10.2 Sample selection

My primary sample consists of the General Population Sample, including the Northern Ireland sample and the Ethnic Minority boost samples. I drop those respondents who completed proxy interviews and all those where relevant information is missing. I restrict the sample to those individuals who are the heads of their respective households and of prime working age (25 to 55), and in full-time employment, i.e. (w_jbstat == 2 & w_jbsemp ==1). I also remove those individuals who are not consistently observed for a continuous spell of at least 3 periods. This leaves a sample of 7,595 individuals with a total of 40,023 observations.

10.3 Definitions of income

For pre-policy labour income, I use monthly gross labour income in the current job (fimnlabgrs_dv) and multiply by 12 to arrive at annual gross labour income. For post-policy income I use monthly net labour income in the current job (fimnlabnet_dv) as well as social benefit income (fimnsben_dv), again annualised. All values are deflated using the annual Consumer Price Deflator for the UK (2015 = 100). I also trim the top and bottom 1% of values of gross labour income in every wave.

10.4 Mincerian regression

In order to partial out the observable components, I run a mincer-type regression of the natural logarithm of gross labour income $\ln(\underline{w}_{i,t})$ on a number of demographic variables:

$$\ln(\underline{w}_{i,t}) = \beta_0 + \sum_{k=1}^{K} \beta_k D_{k,i,t} + \epsilon_{i,t}, \tag{2.27}$$

where $D_{K,i,t}$ contains demographic information about the household:

1. An indicator for the sex of the respondent household member (w_sex_dv).

2. A third-order polynomial of mean household age (calculated from w_age_dv).

3. A dummy for each of the 12 UK government office regions (w_gor_dv).

4. A dummy for the year of the interview (w_intdaty_dv).

I collect the residuals $\epsilon_{i,t}$ and generate a new variable which will serve as the proxy for the pre-policy distribution of labour income targeted by the model: $y_{i,t} = e^{\epsilon_{i,t}}$ which is normalised to have mean 1.

10.5 Human capital proxy

The third wave of Understanding Society was supported by the Economic and Social Research Council (ESRC) and the Department for Business, Innovation, and Skills' Large Facilities Capital Fund, resulting in the addition of a module that included questions about the cognitive and psychological characteristics of adults (16 years and older). I use the results from some of these survey questions to construct a measure of general intelligence, which I use as a proxy for human capital.

I follow Whitley et al. (2016) by constructing a composite measure of cognitive ability by choosing 5 exercises from the cognitive module; these were Numeric Ability, Subtraction, Number Sequence Completion, Word Recall, and Verbal Fluency. I standardise each of these measures individually to have mean 0 and standard deviation of 1 and then perform a Principal Component Analysis on the set of standardised measures.

Table C1 below reports the factor loadings of the variables. All loadings are positive, suggesting that doing better in any type of test is associated with a higher level of general intelligence, which is what one might expect. The loadings are also of similar size, even though numeric ability and the number sequence exercises have slightly higher loadings.

Variable	Loading	Unexplained Variation
Numeric Ability	0.53	0.41
Subtraction Exercise	0.39	0.68
Number Sequence	0.52	0.43
Word Recall	0.38	0.70
Verbal Fluency	0.40	0.65

Table C1: Factor Loadings of PCA on Cognitive Skills

I obtain the first principal component and normalise it to have positive support: h^{UnSoc}. I take the resulting measure to be a proxy for the distribution of human capital which is targeted in the SMM procedure described below.

10.6 Minimum distance procedure

This section contains the details of the calibration procedure to set some of the parameters of the model. The calibration is an application of a Method of Simulated Moments (SMM) estimator (c.f. Guvenen et al. (2014)), or an Indirect Inference (II) procedure (c.f. Gourieroux et al. (1996)). For a given vector of parameters θ, I solve the model and calculate the stationary distribution of households. I then take measures of relevant moments of the distribution of wealth, human capital and income in the model $m(\theta)$ and compare them to relevant data counterparts m^{Data}. Specifically, I minimize the sum of squared percentage deviations of the model-generated moments from their data counterparts:

$$\min_{\theta} \sum \left(\frac{m(\theta) - m^{Data}}{m^{Data}} \right)^2 \tag{2.28}$$

The vector of parameters that minimizes this distance metric is found using a nonlinear solver.

The targets are chosen to represent the distribution of labour earnings, human capital and wealth in the data. Accordingly, the targets can be split into three groups:

1. Gross income targets
   - Mean gross income is given by the mean of the mincerian residual $y_{i,t} = e^{\epsilon_{i,t}}$ described above. The value is 1 by construction.
   - The variance of gross income is given by $Var(y) = 0.288$.
   - The persistence of gross income is the autocorrelation coefficient of a regression of $y_{i,t} = \phi y_{i,t-1} + \varepsilon$.I estimate a value of $\phi = 0.887$ in the sample.
2. Wealth targets
   - The capital-to-income ratio $\frac{K}{Y}$ as a proxy for the aggregate capital stock. I use 2.5 which is a commonly used value.
   - The share of households with no or negative wealth. I use a value of 19% as was estimated from the Wealth and Asset Survey by Angelopoulos et al. (2021).
3. Human capital targets
   - The variance of the human capital measure constructed above: $Var(h^{UnSoc}) = 0.038$.
   - The skewness of human capital is given by $Skew(h^{UnSoc}) = -0.714$.

CHAPTER 3

TO WHAT DEGREE? RECOVERING CHANGES IN THE UK'S GRADUATE SKILL DISTRIBUTION

1 Introduction

It is generally believed that a degree confers (or signals) certain specific skills and abilities. This belief is reflected in the economic literature on the subject of higher education - through the terminology of *human capital* - as well as in the language that Higher Education Institutions use to describe their own function - via terms such as *graduate skills* or *attributes*. Despite this central position, there are few studies that describe what skills, or how much of a specific skill, graduates possess when they finish their university education. At a time when the value of a university degree is coming under increased scrutiny,[1] such quantitative evidence would be very valuable to prospective students, Higher Education Institutions, government departments and employers.

For about a century, a university degree has been seen as a secure route to wealth and professional success. Over the course of mastering their chosen subject, students generally acquire different skills that help them succeed in an increasingly skill-biased labour market (c.f. Goldin & Katz (2009)). Recent studies emphasize the importance

[1] Lately, in the UK and the US, the public discourse has revolved around the value of degrees which - according to some - do not provide relevant skills to students, resulting in high drop-out rates, or poor labour market outcomes for graduates. As a response, the UK government has indicated that funding for such "underperforming" courses might be reduced or cut completely.

of subject of study ("college major") for determining the return to university for a given individual (e.g. Altonji *et al.* (2016), Andrews *et al.* (2022) and Lovenheim & Smith (2022)). Given the wide range of possible subjects students can study, it is natural to assume that differences in subject-specific outcomes are -at least partially - due to the different types of skills that are taught in these courses and how these are rewarded in the labour market. So what types of skills do graduates have? What are the differences between an economics graduate and a medical doctor? And given large changes to the structure of the education system and the wider economy, has the distribution of different skills changed over time?

This paper provides an attempt to quantify the distribution of skills amongst recent university graduates in the UK over the last 25 years. Focussing on two relevant types of skills - mathematical/technical and verbal/organisatorial[2] - I provide estimates of subject-specific skill distributions allowing a quantitative assessment of the skills of a typical graduate as well as the degree of skill inequality between graduates. Taking a long view this paper provides separate estimates for three time periods, covering the period from 1994 until 2019, allowing an evaluation of how the graduate skill distribution has evolved during a period of significant changes in both the primary, secondary and tertiary education sector and the wider labour market.

Over the period under consideration, the UK (and other developed economies) have experienced a rapid expansion of tertiary education participation. Since the passage of the Further and Higher Education Act 1992, university enrolment has roughly doubled to approximately 2 million in recent years. The fact, that these increases were sustained in the face of stark tuition fee increases[3] suggests that a university degree is still seen as a profitable investment by many, but rising graduate earnings inequality and underemployment (c.f. Altonji et al. (2016), Holmes & Mayhew (2016), Lindley & MacIntosh (2015)) cast some doubt on this perception. Generally, it is not clear, how the higher education sector has coped with the rapid expansion of demand and whether a degree still confers the same benefits as it did 30 years ago.

Figure 1 provides an illustration of the evolution of hourly wages for different age categories of graduates. Mean log wages of younger and more experienced graduates comove, even though older workers naturally enjoy higher mean wages. What is of more interest are the diverging trends of wage inequality between the two groups. Since the beginning of the sample period, wage inequality - measured by the variance of logarithms -

[2] Throughout the text I will use "mathematical/technical", "mathematical" and "technical" interchangably. The same applies to "verbal/organisatorial".

[3] The cap on the amount that universities can charge was increased nearly threefold in England in 2012, leading to a large increase in tuition fees with most institutions charging the maximum amount.

is rapidly increasing in the full sample. This finding echoes those of Lindley & MacIntosh (2015), who document rising wage inequality over the period 1994 - 2011. Curiously this pattern seems to be primarily driven by those graduates with more labour market experience since inequality is consistently falling for the group of younger graduates. The pattern persists, even if one focuses on residual wage inequality rather than the raw wages. Wage inequality, particularly focussing on residuals after controlling for observable factors, is indicative of differences in skills and how these are rewarded in the labour market. The observation that residual wage inequality appears to be falling for younger graduates opens the question of how this trend is related to changes in the underlying distribution of skills in the context of large changes to the UK's education system, and corresponding changes in the demand for these skills in the wake of large scale technological change.

Figure 1: Graduate's wages and wage inequality in the UK (1994 - 2019)

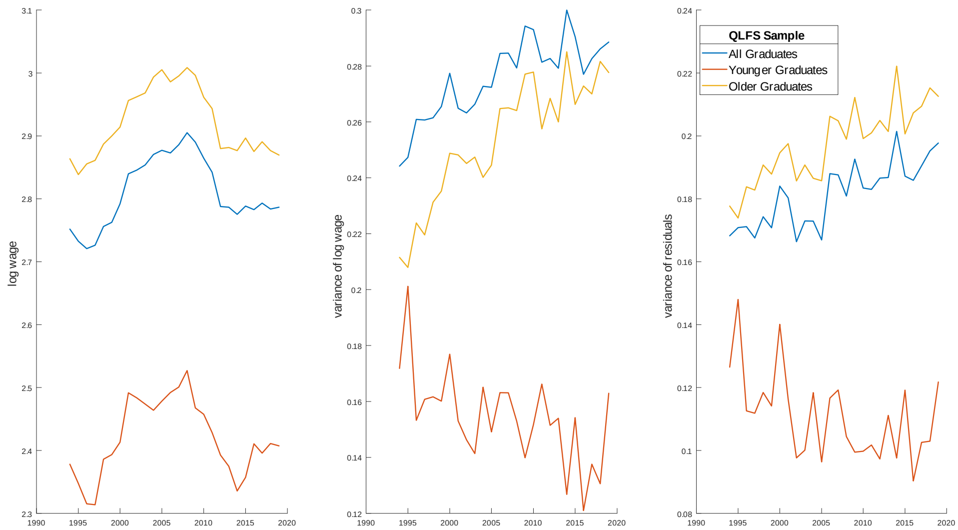

Note: Prime age full time employed graduates. Wages are hourly wages, deflated by 2014 CPI Index.

Residuals are from a mincerian regression of log wages, controlling nonparametrically for age, sex, 1 digit SOC-2000 occupation,

broad degree category (see Section 4), and year and occupation-year fixed effects. Younger graduates are

those between ages 21 and 27 years. Older graduates are those aged between 28 and 55 years.

Source: Quarterly Labour Force Survey

Since direct measures of graduates' skills are lacking in all but the most detailed surveys, specific skills are generally unobservable for an econometrician working with individual-level survey data. In order to address this issue, I take another approach - framing the question as a latent variable problem: skills are unobserved but related to observable choices and labour market outcomes. Hence, by specifying and estimating a corresponding structural economic model, we can make inferences about the unobserved skill endowments of university graduates only using widely available data sources.

To quantify the variation in unobserved skills, I develop a model of occupational choice for university graduates. After graduating, graduates differ with regard to their idiosyncratic endowment of two types of general skills: mathematical/technical and verbal/organisatorial. Skill endowments are modelled as draws from subject-specific multivariate distributions and are thus allowed to vary between as well as within university subjects, capturing important dimensions of heterogeneity: Differences between the distributions capture differences in emphasis due to the specific subject, while each distribution encompasses a further degree of heterogeneity resulting from differences in university quality as well as inherent differences due to individual ability and aptitude conditional on subject choice. Each university subject is therefore characterised by a time-varying, multivariate distribution function, and each graduate by the skill endowment which they have drawn from this distribution.

Just like all graduates have different skills, occupations vary in what value they assign to different skills and hence the match between graduate and occupation matters for realised productivity and wages. The production sector follows the standard approach in the task-skill literature (c.f. Autor et al. (2003), Autor & Handel (2013), Sanders & Taber (2015)), and features a multitude of occupations that differ with respect to how intensely they use each type of skill in production. The combination of a worker's skills and the work task requirements of an occupation determine the worker's occupation-specific productivity. Upon graduation, graduates choose their preferred occupation taking into account their idiosyncratic skill endowment as well as other preferences.

I use a sample of recent university graduates from 1994-2019 together with occupation-level information on work tasks, to structurally estimate the model using simulated maximum likelihood, and recover the parameters of the underlying latent skill distributions for different subjects and time periods. To ensure robustness of the estimates I control for a variety of potential factors that might affect changes in skill demand. I then use the model estimates to analyse changes in the graduate skill distribution and their effects on the labour market outcomes of university graduates over two and a half decades.

The results suggest that since the mid-1990s, there have been substantial changes in the distribution of skills. Over the time period, I find that the median graduate's endowment of effective mathematical skills increased[4], by roughly 140%; effective verbal skills

[4]Beginning in the early 2000s the UK government began to strongly encourage Science, Technology, Engineering and Mathematical education at all levels of education through initiatives like the Science and Innovation Investment Framework. It is therefore not entirely surprising that these policies had a

decreased by around a third. Across 5 major subject categories all but one - Medical and Life Sciences - saw their median mathematical skills rise, most notably Business & Economics which appears to have become increasingly focussed on technical analysis, but also subjects in the Arts & Humanities which saw rises from initially very low levels. Correspondingly most subject categories have lost some of the organisatorial skills that used to be associated with them. Overall the increase in more technical skills has counteracted the fall in the typical graduate's endowment of such "softer" skills, suggesting a reorientation of skill supplies in line with more demand for technical abilities.

To put these changes into perspective: in the mid-90s on average only 27% of a graduate's wage could be attributed to their technical abilities, while organisatorial skills accounted for around 61%. In the period just before the Covid-19 pandemic, these shares have almost reversed, with mathematical skills accounting for around 50% of hourly wages, while verbal skills only contributed around 41%.

In terms of distributions, overall mathematical skills inequality, as measured by the Gini coefficient decreased from 53 to 38 points, verbal skill inequality increased from 24 to 46 points. For STEM and Business & Economics, rising mathematical skills were associated with shrinking within-subject skill inequality, while the converse held for Arts & Humanities and Other degrees.

Together these changes had nontrivial effects on the labour market outcomes of university graduates. In counterfactual simulations I find that in the absence of changes to the subject-specific skill distributions, mean wages would be up to 8% lower whereas wage inequality would be up to 5% larger than what is observed in the data. Additionally, I find that changes in the demographic composition of graduates had only small effects on overall labour market outcomes.

This paper adds to a large, and growing literature on the returns to higher education and specifically to a subset of this literature that investigates the return to specific fields of study (see Altonji et al. (2016), Andrews *et al.* (2022) and Lovenheim & Smith (2022) for extensive surveys). Generally, these studies estimate various latent average treatment effects, whilst trying to address the inherent difficulties caused by the existence of selection effects across dimensions of inherent ability and preference using administrative cutoff rules (see for example Kirkeboen et al. (2016), Hastings et al. (2013)); or try to control for observable factors (Hamermesh & Donald (2008)). This paper retains some uniqueness by making the skills that graduates poses, the key feature of interest, thereby allowing an assessment of the mechanism underlying the returns to different fields of study.

Further, this paper contributes to the literature on how endowments of different types

large impact on the mathematical and analytical skills of graduates from the mid-2000s onwards.

of skills affect labour market outcomes of graduates in an environment where occupations have differentiated skill requirements. These papers tend to focus on the dichotomy between more general (transferable) and more specific skills leading to differences in the risk-return profiles between general and specialised degree subjects (c.f. Leighton & Speer (2020), Onozuka (2019)). Of particular importance here is the paper by Kinsler & Pavan (2015), which estimates a structural model where students acquire mathematical and verbal skills the return of which differ according to their occupation. The modelling approach taken in their paper is necessarily different from my own, but they are related in spirit.

Finally, this paper complements other attempts at eliciting the skill content of different university degrees. Altonji et al. (2014), create measures of the task content of different subject by mapping task measures from the Dictionary of Occupational Titles to graduate's occupation choices. Similarly, Hemelt et al. (2021) collect information from online job postings, to associate desired skills with different degree subjects. My paper differs in so far as that it uses both occupation choice and wage information for estimation as well as allowing for substantial within-subject skill heterogeneity. However, it shares the former's conception of a university degree as a bundle of multidimensional skills that are related to different tasks.

To the best of my knowledge, this paper is the first attempt at trying to find quantitative evidence for the actual distribution of skills of university graduates in the UK. The results suggest that skill heterogeneity plays a large role in explaining the changes in the labour market outcomes of university graduates. Graduates differ in their skill endowments in accordance with the subject that they choose to study and beyond. Furthermore, the distribution of graduate's skills is changing over time meaning that graduates today look very different from those 10-25 years ago. This finding has important implications for educational and more general economic policy going forward.

The rest of the paper is structured as follows: section 2 presents the economic model of wage setting and occupational choice; section 3 presents the econometric strategy, used to estimate the parameters of interest; section 4 presents the data sources used in the analysis; section 5 highlights the estimation procedure; section 6 covers the results; section 7 presents counterfactual experiments and section 8 concludes.

2 Model

In this section, I present an economic model of occupation choice and wage determination for recent university graduates in order to recover the skills supplied by university degrees.

The economic environment in this model closely follows the literature on estimating task returns (c.f. Autor & Handel (2013), Roys & Taber (2016)). Whereas for occupation choice, I follow the methodological approach of the multinominal choice literature, where it is common to estimate unobserved parameters from the observed choices of individuals. In particular, I will refer to the class of mixed logit models which seem to be particularly relevant in this context (see Train (2009), Chapter 6). For expositional simplicity, both parts are presented separately, before being combined in the next section.

2.1 Wage determination

A worker's multidimensional skill set is summarized by a K dimensional vector $s_i = \{s_{i1}, s_{i2}, .., s_{iK}\}$ where each element $s_{ik} \geq 0$ describes how effective worker i is at performing task k.

On the firm side, the labour market consists of a large number of competitive firms of different types (henceforward referred to as "*occupations*") that use the different skills supplied to them in different proportions. Specifically, every occupation $o \in O$ has an associated vector $\lambda_o = \{\lambda_{o1}, \lambda_{o2}, ...\lambda_{oK}\}$ where each element $\lambda_{ok} \geq 0$ summarizes the productivity of task k in occupation o.

A worker's human capital therefore depends on her skill set as well as the task-productivity vector of her chosen occupation. Specifically, the human capital of worker i in occupation o is defined as:[5]

$$h_{io} = e^{\sum_{k=1}^{K} \lambda_{ok} s_{ik}} \tag{3.1}$$

Denote the aggregate amount of human capital in occupation o as

$$H_o = \int_{i \in o} h_{io} d(i). \tag{3.2}$$

Finally, output is produced by an aggregate production function:

$$Y = F(H_1, ..., H_O) \tag{3.3}$$

The marginal product of worker i in occupation o is:

$$\frac{\partial Y}{\partial h_{io}} = \frac{\partial F}{\partial H_o} \frac{\partial H_o}{\partial h_{io}} = \frac{\partial F}{\partial H_o} e^{\sum_{k=1}^{K} \lambda_{ok} s_{ik}} \tag{3.4}$$

[5] For this exposition I am going to ignore any other factors that might influence productivity such as worker specific characteristics. Including these is a trivial extension of the model.

Denote $\frac{\partial F}{\partial H_o} = e^{\eta_o}$,[6] and assume that firms pay workers their marginal product, then the log wage of worker i in occupation o can be written as:

$$w_{io} = \eta_o + \sum_{k=1}^{K} \lambda_{ok} s_{ik} \tag{3.5}$$

This setup is fairly standard in the literature on tasks and skills (c.f. Autor & Handel (2013), Roys & Taber (2016)).

Generally, economies of the type described above are characterized by the sorting of workers according to comparative advantage (see Roy (1951)). This self-selection of workers into different occupations according to their different abilities poses the main obstacle that is faced by the literature that is concerned with estimating "task prices" (i.e. the set λ). Since there will be a positive correlation between an occupation's task prices λ_o and the skills supplied by workers selecting into this occupation, simply running an OLS regression on equation (3.5) will not do the trick (see Autor (2013)).

In this paper, however, rather than being harmful, self-selection is actually helpful as it allows us to make inferences from a worker's observed occupation to her unobserved skill set. In order to enable this inference, we will first have to model the discrete choice behaviour of the worker.

2.2 Occupational choice

Workers observe their skills, and all potentially relevant characteristics of an occupation and pick whichever occupation provides them with the highest valuation in terms of utility. In this case, suppose that every graduate can observe the set O of all available occupations and attach a personal valuation V_{io} to each of these options. Accordingly, a worker i solves the following (static) occupational choice problem:

$$V_i = \max_{o \in O} \{V_{io}\} \tag{3.6}$$

Under these circumstances the individual's occupation choice o_i^* will refer to the best available option:

$$o_i^* = \arg\max\{V_{io}\} \tag{3.7}$$

In the following, I will make some assumptions about the different parts affecting the worker's utility V_{io} which allows me to estimate the unobserved characteristics that we

[6] Depending on your preferences, you might want to interpret η_o as an occupation-specific demand component, or an occupation fixed effect.

are interested in. Let us assume that the utility derived from the occupation is linear in the log wage,[7] leading to the following relationship:

$$V_{io} = w_{io} + \varepsilon_{io} \tag{3.8}$$

where o is one of the available occupations, w_{io} is the log wage earned by i in occupation o and ε_{io} is an individual-occupation-specific preference shock that is **i.i.d.** across all agents and all occupations.[8] Importantly, the value of V_{io} is perfectly observed by the economic agent, while only o_i^* is observed by the econometrician.

The random component ε_{io} is random in an idiosyncratic sense. Two workers with the same deterministic wage may have different preferences over the set of occupations. This differentiation in choice behaviour is important since otherwise, the utility-maximizing choice would be the same for every worker, leading to unrealistic predictions. Furthermore, the introduction of this random term allows us to capture other factors that influence occupation choice besides the desire to maximize wages, such as other preferences or frictions in the labour market.

3 Econometric Strategy

The econometric strategy combines the empirical content of the two parts of the economic model described above. The key ingredient is that both, a worker's occupation choice and her realized wage are informative about her skill set, provided that we also have some information about the occupation task vector λ.

3.1 A mixed logit model of occupational choice

Let us recall the problem our graduate is facing. She knows her own skill set s_i, as well as the task vectors of all occupations λ, as well as the occupation-specific parameters η, and therefore perfectly knows her log wage in every occupation o: $w_{io} = \eta_o + \sum_{k=1}^{K} \lambda_{ok} s_{ik}$.

She also perfectly knows her preferences over the non-pecuniary aspects of each occupation ε_i, and is therefore able to assign to each occupation a personal valuation

[7] This is likely to be the case for an economic agent with a suitably defined utility function (e.g. logarithmic), who is borrowing constrained. I believe it reasonable to assume that this situation applies to the sample population studied in this paper.

[8] Generally, models of this form are known as "random utility models" (RUM), since the worker's valuation of the different options V_{io} can be broken up into a "deterministic" part, w_{io} and a "random" part, ε_{io}.

$V_{io} = w_{io} + \varepsilon_{io}$. Finally, given this valuation the graduate chooses her preferred occupation: $o_i^* = \arg\max\{V_{io}\}$.

Making the standard assumption that her idiosyncratic occupation preference shocks ε_i are distributed i.i.d. Type I Extreme Value, we can express the conditional choice probability of her chosen occupation o_i^* as:

$$\Pr(o_i^*|s_i) = \frac{e^{w_{io^*}}}{\sum_{o=1}^{O} e^{w_{io}}} = \frac{e^{\eta_{o^*} + \sum_{k=1}^{K} \lambda_{o^*k} s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o + \sum_{k=1}^{K} \lambda_{ok} s_{ik}}} \tag{3.9}$$

Assuming that s_i was drawn from a parametric distribution, then it is possible to identify and estimate the parameters of this distribution.[9]

I assume that each skill vector is drawn from a multivariate log-normal distribution with mean μ and variance-covariance matrix Σ:

$$\log(s_i) \sim MVN(\mu, \Sigma). \tag{3.10}$$

The log-normal is a convenient choice here, as it ensures strictly positive support for the skill set s, which seems like a reasonable choice for our purposes.

Using this assumption, we can derive the unconditional choice probability by integrating over the distribution of s:

$$\Pr(o_i^*) = \int \Pr(o_i^*|s_i) f(s) d(s) = \int \frac{e^{\eta_{o^*} + \sum_{k=1}^{K} \lambda_{o^*k} s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o + \sum_{k=1}^{K} \lambda_{ok} s_{ik}}} f(s) d(s) \tag{3.11}$$

Standard results (c.f. McFadden & Train (2000)) guarantee, that we can use the unconditional choice probability in (3.11) to get consistent estimates for η, μ and Σ, using simulated maximum likelihood.[10]

However, the model is not complete yet. As of yet there is nothing distinguishing our "skill" interpretation of s from a "taste" interpretation. Indeed, strictly speaking we would have to provide a location normalization for one of our parameters, in order to fix their relative values. In the following I will use the observed wage to address the last two points.

[9] Assuming a parametric distribution for s turns this into a mixed logit model (see Train (2009), Chapter 6), where we are effectively treating skills as random taste parameters over the different tasks. The mixed logit is an extremely flexible choice model that can indeed approximate any random utility model (c.f. McFadden & Train (2000)). Most interesting for researchers is that it naturally generates correlations in choice behaviour across similar alternatives. For example, a worker with a particularly large value of some skill is going to prefer all occupations that use this skill with great intensity.

[10] There is no closed form solution for this integral, but the integration step can be performed via simulation.

3.2 Adding wage information

So far the model has already made use of the wage setting equation (3.5), but for any draw of s_i a worker's modelled wage w_i differs from the worker's realized (observed) wage w_i^{obs}, due to the presence of other factors such as individual effort and luck. I capture these elements by adding an additional disturbance term to the wage equation:

$$w_i^{obs} = w_i + \upsilon_i \tag{3.12}$$

where υ_i is a random, mean zero disturbance, **independent** of the worker's occupation choice:

$$\nu_i \sim N(0, \phi^2).$$

As such υ_i does not impact the graduate's occupation choice, as can be easily derived from the analytic form of the occupation choice probabilities. To see this, add υ_i to all potential wage outcomes w_{io}, then the logit formula implies:

$$\Pr(o_i^*|s_i) = \frac{e^{w_{io^*}+\upsilon_i}}{\sum_{o=1}^{O} e^{w_{io}+\upsilon_i}} = \frac{e^{\upsilon_i} e^{w_{io^*}}}{e^{\upsilon_i}\sum_{o=1}^{O} e^{w_{io}}} = \frac{e^{w_{io^*}}}{\sum_{o=1}^{O} e^{w_{io}}} \tag{3.13}$$

Hence, as long as υ_i does not vary across different potential occupations, the choice probabilities remain unaffected.[11]

Given a specific occupation choice o_i^*, and skill-set s_i, we can calculate the size of ν_i:

$$\nu_i = w_i^{obs} - \left[\eta_{o^*} + \sum_{k=1}^{K} \lambda_{o^*k} s_{ik}\right] \tag{3.14}$$

Thinking in terms of the estimation strategy, ν_i provides a measure, of how far the wage implied by the model parameters, is from a worker's actual observed wage. Jumping ahead a little, we should expect the *true* model to minimize this distance.

As ν_i is normally distributed, we have a closed form expression for the conditional probability of observing the observed wage, conditional on a certain skill set s_i and occupation choice o_i^*:

$$\Pr(w_i^{obs}|s_i, o_i^*) = \frac{e^{(-\frac{\nu_i^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}} \tag{3.15}$$

Ultimately, I am interested in finding the set of parameters, that maximizes, the

[11] The same is also true if ν_i varies across occupations, but is unanticipated by the graduate at the time she chooses her occupation.

unconditional joint probability that a worker chooses the occupation that she is observed choosing and that she earns the wage that she is observed earning: $\Pr(o_i^*, w_i^{obs})$.

To find the correct expression, we first rewrite $\Pr(o_i^*, w_i^{obs})$ as

$$\Pr(o_i^*, w_i^{obs}) = \int \Pr(o_i^*, w_i^{obs}|s_i) f(s) d(s) \tag{3.16}$$

using the law of conditional probabilities to rewrite:

$$\frac{\Pr(o_i^*, w_i^{obs}|s_i)}{\Pr(o_i^*|s_i)} = \Pr(w_i^{obs}|s_i, o_i^*) \tag{3.17}$$

$$\Pr(o_i^*, w_i^{obs}|s_i) = \Pr(o_i^*|s_i) * \Pr(w_i^{obs}|s_i, o_i^*) \tag{3.18}$$

Plugging the expression back in gives us:

$$\Pr(o_i^*, w_i^{obs}) = \int \Pr(o_i^*|s_i) \Pr(w_i^{obs}|s_i, o_i^*) f(s) d(s). \tag{3.19}$$

From (3.9) we know that:

$$\Pr(o_i^*|s_i) = \frac{e^{\eta_{o^*} + \sum_{k=1}^{K} \lambda_{o^* k} s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o + \sum_{k=1}^{K} \lambda_{ok} s_{ik}}} \tag{3.20}$$

and hence we can combine to write:

$$\Pr(o_i^*, w_i^{obs}|s_i) = \left(\frac{e^{\eta_{o^*} + \sum_{k=1}^{K} \lambda_{o^* k} s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o + \sum_{k=1}^{K} \lambda_{ok} s_{ik}}} \right) \left(\frac{e^{(-\frac{\nu_i^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}} \right) \tag{3.21}$$

finally, integrating over the distribution of s leads to the unconditional joint probability:

$$\Pr(o_i^*, w_i^{obs}) = \int \left\{ \left(\frac{e^{\eta_{o^*} + \sum_{k=1}^{K} \lambda_{o^* k} s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o + \sum_{k=1}^{K} \lambda_{ok} s_{ik}}} \right) \left(\frac{e^{(-\frac{\nu_i^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}} \right) \right\} f(s) d(s) \tag{3.22}$$

In the appendix I describe a complete algorithm that can be used to estimate the parameters of interest from this model, using the likelihood function implied by (3.22).

3.3 Model extensions

Other demographic characteristics

For an empirical application, it is necessary to control for a number of observable characteristics, as well as circumstantial factors. However, it is trivial to extend the model to include factors other than the skills considered above. To show this, I extend the log wage equation below:[12]

$$w_{io} = \eta_o + \sum_K \lambda_{ok} s_{ik} + \beta x_i \tag{3.23}$$

where x_i is a vector of observable characteristics (gender, labour market experience, etc.), β is a vector of coefficients. Clearly, this equation can be inserted into the likelihood function (3.22), and β can be estimated as part of an extended parameter vector θ. Further, as long as neither x_i, nor β, vary across occupations (i.e. the model does not include for example either occupation-specific experience (occupational tenure) among the observables, nor occupation-specific coefficients in β), the additional terms do not have any impact on the occupational choice probabilities, and can therefore be ignored in the first part of the likelihood calculation.

Systematic occupation preferences

In order to make the model more realistic, I also include systematic non-pecuniary aspects of occupations that might affect the graduate's choice. Specifically, I augment the graduates expected payoff from choosing occupation o by a non-random occupation preference term ω_o, which is constant for all graduates and represents the (dis-) utility of working in a specific occupation. The augmented occupational valuation equation thus reads as follows:

$$V_{io} = w_{io} + \omega_o + \varepsilon_{io} \tag{3.24}$$

Like β, ω can be estimated as part of the extended parameter vector θ. Since ω does only affect the occupational choice probabilities, it can be ignored in the wage equation part of the likelihood function.

[12] Naturally, this can be understood as an extension of the human capital equation specified above.

4 Data

4.1 Graduates

The main data source used in this paper is the Quarterly Labour Force Survey (QLFS) over the period 1994-2019, which I split into three periods: 1994-2002, 2003-2011 and 2012-2019. Since 1994 the QLFS has included reasonably fine-grained information on the subject of an individual's first university degree (see Lindley & MacIntosh (2015) for more details).[13] Furthermore, the QLFS also contains information on an individual's current occupation, usual hourly pay and some other demographic covariates.

I restrict the sample to full-time working graduates between the ages of 21 and 27, who have not graduated more than 2 years before I observe them in the sample.[14] This age restriction is put in place to make sure that we capture those graduates who are "fresh" out of university so that their skill set most accurately reflects their post-university endowment. A small age bracket also reduces contamination by other factors such as age and experience effects as well as on-the-job skill accumulation.

For each graduate in my sample, I collect wages measured as usual gross hourly pay, deflated by the CPI;[15] their current occupation as classified by the 1-Digit SOC Occupational classification schedule; Gender; Subject of First degree; and years since graduation, which I use as a proxy for labour market experience.

I split the sample into 5 groups according to broadly defined subject degree categories: 1. Medical and Life Sciences (including Biology & Agriculture); 2. Science, Engineering, Technology & Mathematics; 3. Business Management and Economics; 4 Arts & Humanities; 5. Other Degrees. In order to avoid complications I drop all those who hold any advanced degrees beyond the undergraduate level.[16] A table summarizing the resulting sample can be found in Appendix A.

The total sample includes 10,669 graduates, with around 3,500 individuals in each time period. Over the 25-year period, the most significant change is shown by the composition of subjects represented among graduates. STEM, Business & Economics and Arts &

[13] I ignore those who have more than one degree, or any further or higher degrees. Postgraduate qualifications take on a more significant role over time, as a higher percentage of graduates pursue these degrees. However, a large fraction of graduates pursue a postgraduate qualification in a subject different from their first degree, making it difficult to assign the skills they exhibit to their first or later degree. I decided to exclude postgraduates in order to keep the relationship between degree subject and skills as clean as possible. However, including higher degree holders without any further changes to the model does not change the quantitative predictions of the model with regard to skill supply or inequality (see Appendix).

[14] Typically in the UK students finish high school at 18 and enter 3-year University Courses.

[15] I also trim the top and bottom 1% of wage values to remove nonsensical values.

[16] Results are robust to including postgraduates. See Appendix for details.

	Mean hourly wage			Gini hourly wage			HHI index of occupation concentration			Share of nontypical occupations		
	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	10.61	11.54	10.45	0.17	0.17	0.16	0.29	0.25	0.25	12.15	18.93	20.43
STEM	10.81	11.6	11.4	0.18	0.17	0.16	0.25	0.26	0.26	12.72	16.45	15.52
Business & Economics	10.18	10.87	10.41	0.18	0.17	0.17	0.22	0.22	0.21	4.22	7.44	7.38
Arts & Humanities	9.2	9.66	9.3	0.19	0.17	0.16	0.19	0.19	0.19	10.4	14.36	16.38
Other Degrees	10.28	10.86	9.86	0.18	0.18	0.18	0.25	0.25	0.17	9.38	11.48	20.79
All	10.24	10.9	10.32	0.18	0.18	0.17	0.21	0.2	0.2	8.42	10.93	13.78

Note: Wages are CPI deflated (2014 = 100).

Table 1: Summary statistics of QLFS sample - Labour Market Outcomes

Humanities lose a proportion of their graduates, while Medical & Life Sciences and the category of other degrees gain in relative popularity. The differential growth of different degree subjects might also be reasonably linked to broader demographic factors, such as for example increasing participation of women in higher education as well as in the labour market. Although the share of women increased by around 4 percentage points, the relative share of women among the different subjects has remained relatively stable. Women are most underrepresented among STEM subjects - making up around a quarter to a third of all graduates, while they make up around 70% of Medical and Life Sciences.

The labour market outcomes of these recent graduates are summarised in Table 1. The average hourly real wage in the sample is relatively stable at just over £10/hour, even though it comes close to £11/hour during 2003-2011. Across subjects STEM graduates consistently earn the highest average wage, closely followed by Medical & Life Sciences and Business & Economics graduates. Arts & Humanities graduates tend to have the lowest average wages. Apart from these between-subject differences, there is a large variation of within-subject wage inequality with within-subject gini coefficients of around 18 gini points which is comparable to the overall gini coefficient. Across time there appears to be a slight decline in wage inequality: between the period 1994-2002 and 2012-2019, the overall gini coefficient of the hourly wage declines somewhat from 18 to 17 gini points.

To measure occupational outcomes, I also report the Herfindahl-Hirschman Index of occupational concentration as well as my own measure of the "share of non-typical occupations", which I define as the share of graduates that work in an occupation that used to have a share of <5% of graduates in the first period. Occupational concentration appears to somewhat decrease in Life Sciences and Other graduates, while it increases amongst STEM graduates. Overall the effect is a small reduction in occupational concentration. However, simply looking at the HHI disguises an important trend that becomes evident when we look at the share of graduates entering occupations that they would not have traditionally entered.[17] Here the share grows by approximately 60% overall, although there is some variation across different subjects. This broad trend might provide some evidence as to the theory of increasing underemployment of university graduates.

[17] These occupations typically include Skilled Trades, Service and Sales or Elementary Occupations.

4.2 Occupations & tasks

The one-digit SOC 2000 schedule provides me with 9 occupation groups. For the task dimension, I choose two broad groupings: 1. Mathematical/Technical Tasks; 2. Verbal/Organisatorial Tasks. I choose these groups since I believe that these kinds of tasks are of particular relevance to university graduates.[18]

To obtain an estimate of the occupation task requirements, I use four waves of the UK Skills and Employment Survey (SES). The years of these surveys, 1997 & 2001, 2006 & 2012 and 2017 map neatly into our sample periods. Since the beginning of the task literature, there have been many different approaches that try to approximate the task requirement vector λ using survey data (c.f. Autor (2013), Autor et al. (2003), Autor & Handel (2013), Rohrbach-Schmidt & Tieman (2013)). Here I follow the approach of Bisello (2013) who also works with the SES. In this survey, respondents answer questions related to their job and score the importance of performing certain tasks on a Likert scale. I group some of these questions into the three task dimensions and perform a dimensionality reduction using principal component analysis. I then scale the obtained values between 0 and 1 and average across occupations. The resulting task vectors are summarized in Figure 2.[19]

Without going into too much detail, I would like to highlight the changes to task requirements over the time period. These have been almost exclusively positive across the two tasks between the first and the second period with a slight reversal between periods two and three. Such changes would be very much in line with any explanation emphasizing the increasing role of Cognitive and Non-Cognitive skills as a result of increased Information and Communications Technology (ICT) usage (c.f. Acemoglu & Autor (2011)). For our purposes, it appears important to account for the changing task requirements over time, as they generate important variation that is useful to identify the changing parameters of the skill distributions.

[18]For example I do not include manual or routine tasks, as I do not believe that these are particularly interesting in the context of higher education.

[19]More details are provided in the Appendix.

Figure 2: Task Weights by 1-Digit SOC 2000 Occupation

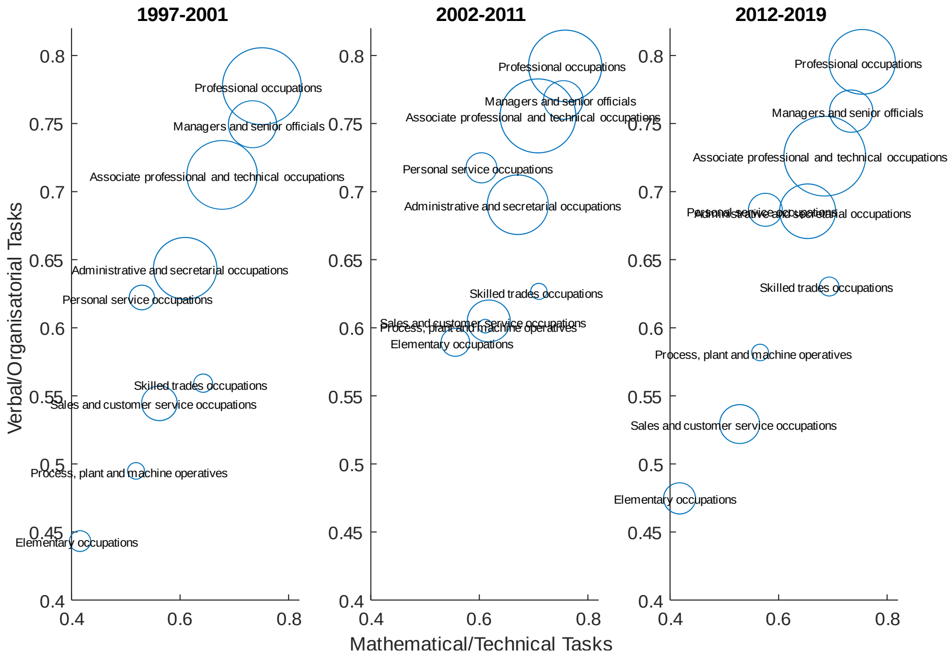

Note: Circle size proportional to employment share in QLFS sample.

5 Estimation

Let's recall that we are interested in estimating the parameters of the subject-specific graduate skill distribution, which had been specified as: $\log(s_i) \sim MVN(\mu_t, \Sigma_t)$.

I want to recover changes in the skill distribution over time, so both μ_t and Σ_t are specified as time-varying. I have specified two task dimensions and correspondingly the skill distribution also has two dimensions $k = \{1, 2\}$. Furthermore, there are 5 degree subjects, $m = \{1, .., 5\}$ and three time periods $t = \{1, 2, 3\}$, leading to $m * t = 15$, subject-period specific multivariate skill distributions. For the covariance structure, I assume that skills are uncorrelated within each subject-period distribution:

$$\Sigma_{mt} = \begin{bmatrix} \sigma^2_{1,mt} & 0 \\ 0 & \sigma^2_{2,mt} \end{bmatrix} \tag{3.25}$$

Note that this doesn't imply that skills are uncorrelated at the population level. If a certain subject-period combination generates high values of two different skills, it will indeed look like there exists a positive correlation between these two skills. It is only assumed that there is no correlation within each subject.

Occupation fixed effects (η_{ot}) are also allowed to vary between the two periods, in order to capture structural changes in the demand for their output. Across the three periods, the sample spans 26 years, and I allow for year-specific aggregate conditions in the labour market, by including year fixed effects. I also include a linear term for experience and gender, both of which are allowed to vary across periods.

To summarize, we have to estimate 60 parameters (μ_{kmt} & σ^2_{kmt}) for the 30 different lognormal distributions, 24 for the occupation fixed effects η_{ot}, 24 for the occupation-specific preference terms ω_{ot} , 23 year fixed effects and 3 each for gender and experience controls - a total of 137 parameters.

Setting ϕ^2, i.e. the variance of the measurement error, is a difficult task in this model, that requires some additional steps. The error term υ_i does not only capture traditional measurement error but also any other productivity differentials that materialize over the course of the graduate's early career, such as health episodes or promotions. The standard approach to setting ϕ^2 would be to run a regression of wages on a number of observables and use the variance of the residuals as an estimate. For this model, this requires controls for mathematical and verbal abilities. Luckily, I can resort to an auxiliary data set (Understanding Society, Wave 3), providing me with an estimate for $\phi = 0.17$.[20] The details of the estimation algorithm are provided in the Appendix.

6 Results

This section presents the results of the estimated model. I first present an evaluation of the model's fit to the data. Then I discuss the changes in the underlying unobserved skill distributions and their implications for the observed wage dynamics. Finally, I present the results of some counterfactual experiments in the next section.

6.1 Model fit & validation

I evaluate the ability of the model to capture both the occupation choices of graduates as well as their wage outcomes. For this purpose, I simulate a random, representative sample of $100,000$ graduates in each time period. Figure 3 below shows the histogram of the hourly wage across all time periods, while Figure 4 highlights the model fit with respect to the occupation choices of graduates in each time period. Across both dimensions the model tracks the data very well, capturing both the shape of the wage and occupation distributions and tracking their changes across time.

[20]For more details, see Appendix. Results are robust to increasing or decreasing ϕ by 10%.

Figure 3: Histogram of Hourly Wages.

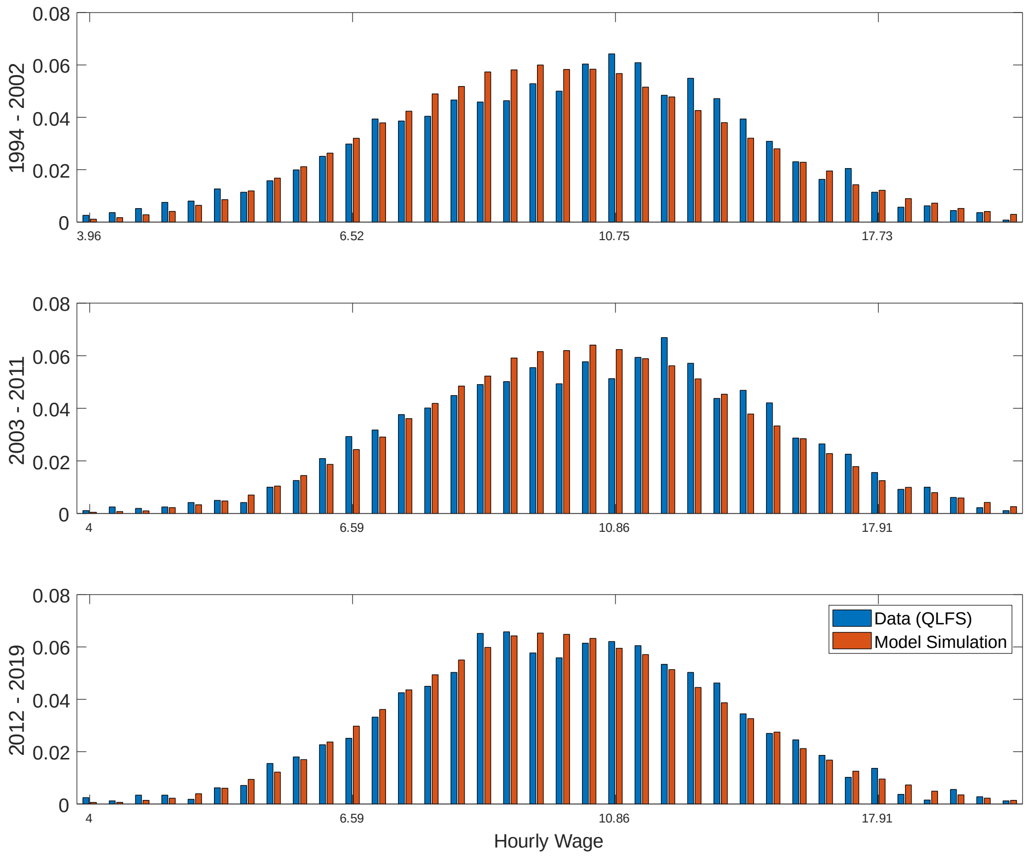

Notes: Histogram of hourly wages. QLFS Data and Simulation. Wages in the data are deflated by the 2014 CPI Index.

Figure 4: Occupation Distribution - QLFS Data and Model Simulation.

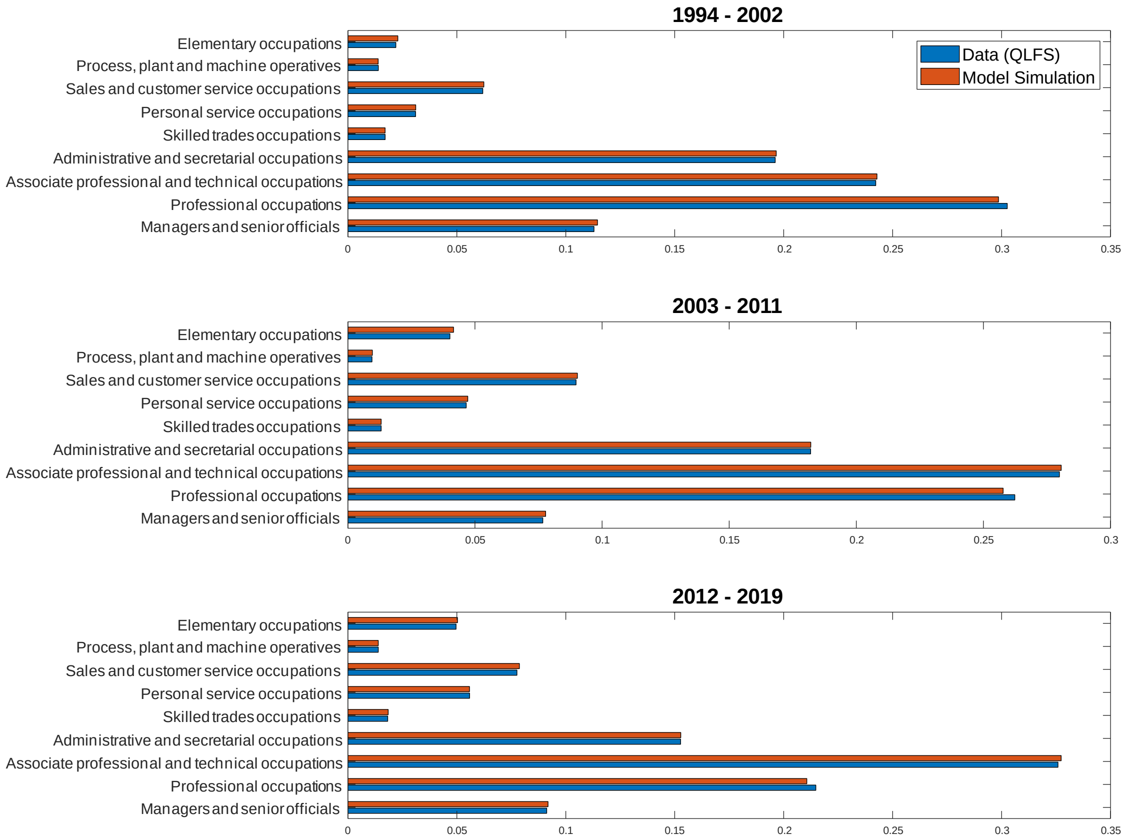

Notes: 1-Digit SOC 2000.

To complement the visual with some statistical evidence the table below compares the model and the data with respect to the mean wage, the Gini coefficient of the hourly wage and the occupation concentration index, as well as the share of non-typical occupations. The model fit is quite good, with the overall model predictions matching their empirical counterparts closely (usually within 1% margin of error).

At the subject level, the predictions perform slightly worse with the wage indicators (mean and gini). The subject-specific means are actually captured quite well to within a maximum deviation of around 5%. Within-subject gini coefficients are generally slightly overpredicted to a maximum of around 8% above the values observed in the data. Overall, these deviations are not too concerning and I argue that the model still does a good job of capturing between and within-subject differences in wages. In terms of subject-specific occupational outcomes, the model has slightly more difficulty matching the data indicators for both the HHI occupational concentration index and the share of non-typical occupations. This is likely due to the fact that the model differentiates occupations only by the task weights and does not include any subject-occupation-specific skills or human capital.[21]

[21] For example, if the occupations of medical doctor and engineer have similar task weight vectors, then the model would predict that a graduate should be approximately indifferent in choosing either occupation, notwithstanding that in reality there are obvious additional factors that determine whether one chooses to become one rather than the other.

Mean Wage									
	1994 - 2002			2003 - 2011			2012 - 2019		
	Data	Model	Difference (%)	Data	Model	Difference (%)	Data	Model	Difference (%)
Medical and Life Sciences	10.61	10.20	-3.82	11.54	11.14	-3.48	10.45	10.26	-1.87
STEM	10.81	10.51	-2.76	11.60	11.09	-4.42	11.40	10.76	-5.59
Business & Economics	10.18	10.64	4.55	10.87	11.26	3.55	10.41	10.60	1.89
Arts & Humanities	9.20	9.62	4.62	9.66	10.19	5.52	9.30	9.60	3.24
Other Degrees	10.28	10.15	-1.30	10.86	10.49	-3.41	9.86	10.09	2.34
All Degrees	10.24	10.25	0.09	10.90	10.84	-0.52	10.32	10.27	-0.45

Gini Wage									
	1994 - 2002			2003 - 2011			2012 - 2019		
	Data	Model	Difference (%)	Data	Model	Difference (%)	Data	Model	Difference (%)
Medical and Life Sciences	0.171	0.185	8.328	0.171	0.177	3.351	0.160	0.161	0.452
STEM	0.180	0.190	6.032	0.169	0.177	4.991	0.155	0.167	7.704
Business & Economics	0.183	0.192	5.136	0.169	0.178	5.053	0.170	0.177	4.328
Arts & Humanities	0.188	0.190	1.144	0.167	0.172	2.844	0.157	0.168	6.759
Other Degrees	0.179	0.182	1.300	0.177	0.175	-0.958	0.176	0.178	1.490
All Degrees	0.184	0.189	3.137	0.176	0.177	0.879	0.169	0.171	1.476

HHI of Occupation Concentration									
	1994 - 2002			2003 - 2011			2012 - 2019		
	Data	Model	Difference (%)	Data	Model	Difference (%)	Data	Model	Difference (%)
Medical and Life Sciences	0.29	0.21	-29.36	0.25	0.20	-21.26	0.25	0.19	-23.60
STEM	0.25	0.21	-18.88	0.26	0.19	-24.80	0.26	0.20	-25.55
Business & Economics	0.22	0.21	-6.99	0.22	0.20	-11.32	0.21	0.20	-6.68
Arts & Humanities	0.19	0.20	8.84	0.19	0.20	5.59	0.19	0.20	0.27
Other Degrees	0.25	0.20	-18.51	0.25	0.20	-20.35	0.17	0.20	15.11
All Degrees	0.21	0.21	-0.71	0.20	0.20	-0.77	0.20	0.20	-0.22

Share of nontypical Occupations									
	1994 - 2002			2003 - 2011			2012 - 2019		
	Data	Model	Difference (%)	Data	Model	Difference (%)	Data	Model	Difference (%)
Medical and Life Sciences	12.15	14.45	18.87	18.93	20.16	6.49	20.43	21.93	7.34
STEM	12.72	14.72	15.66	16.45	20.18	22.71	15.52	21.60	39.18
Business & Economics	4.22	8.61	104.01	7.44	10.69	43.70	7.38	12.98	75.77
Arts & Humanities	10.40	8.92	-14.19	14.36	11.50	-19.91	16.38	13.83	-15.56
Other Degrees	9.38	11.41	21.71	11.48	15.18	32.17	20.79	16.51	-20.58
All Degrees	8.42	8.51	1.03	10.93	11.15	2.02	13.78	13.87	0.61

Note: Wages are CPI deflated (2014 = 100).

Table 2: Model Fit

6.2 Skill results

Unfortunately, the shape parameters of a log-normal distribution are not particularly intuitive, so I have presented the effective median and mean skill levels and changes in the next table. Since the lognormal is not symmetric there is a difference between the median and the mean outcome. Overall the trends and results are very similar if we take the mean instead (reported in the same table). For this part, I will focus the discussion on the median, as it provides a convenient interpretation of the skills of a "typical" graduate.

Looking at these we can first confirm some of our initial expectations, for example, STEM subjects seem to endow their students with more mathematical skills, whilst somewhat lacking in the verbal department. Further observations of this sort should convince us that the estimation is actually picking up some *real* differences between subjects.

Looking at the average *effective* skill levels, we see different trends across both dimensions: Mathematical skills increased by around 120% between 1994-2002 and 2003-2011 and by just over 140% between 1994-2002 and 2012-2019. The increase appears to be - at least in part - driven by increases in the level of technical skills of graduates from subjects that had very low levels of mathematical/technical skills at the beginning of the sample period. Particularly notable is the large increase in the median level of these skills by Business & Economics graduates, but Arts & Humanities and Other Degree graduates also show increasing levels of these skills over the time period. The only exception to this pattern are Medical & Life Science graduates where technical skills become less prevalent. It is tempting to suggest that increased demand for technical skills in the labour market has provided the incentives for these observed patterns.

The increase in mathematical/technical skills is counteracted by a large decrease in verbal/organistatorial skills, which fall by roughly 33% between periods one and two. This decrease remains roughly constant between the second and third periods. This trend is reflected fairly uniformly across subject areas with the exception of Medical & Life Science which experiences the opposite trajectory. These results sugges a changing skill composition amongst graduates, with a rising emphasis on hard cognitive skills that took hold, particularly around the turn of the millennium. The observed increase in technical skills, together with the fall in organisatorial skills, suggests a change in the skill composition of graduates, not necessarily an overall fall in skill levels. To assess the net effect, I sum over both skills, to get an estimate of the overall skill level of graduates. The results suggest a modest increase in median skill levels of around 6% across the sample period. Overall, I believe the lesson to be learned here is that graduate quality has not deteriorated in the wake of the higher education expansion. This finding is consistent with Blundell et al. (2016) who suggest that a significant decline in unobserved ability of graduates was inconsistent with observed wage and employment movements.

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Median Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.44	0.00	0.31	1.43	2.93	2.66	2.88	2.93	2.98
STEM	2.03	3.01	3.15	0.87	0.02	0.00	2.92	3.04	3.15
Business & Economics	0.60	1.59	3.13	2.28	1.40	0.00	2.89	2.99	3.14
Arts & Humanities	0.00	1.55	1.05	2.71	1.31	1.84	2.71	2.87	2.90
Other Degrees	0.00	0.39	1.65	2.82	2.46	1.34	2.82	2.86	2.99
All Degrees	0.66	1.46	1.58	2.14	1.42	1.40	2.85	2.94	3.03

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Mean Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.46	0.00	0.32	1.44	2.94	2.67	2.90	2.94	2.98
STEM	2.05	3.02	3.16	0.89	0.03	0.00	2.94	3.05	3.16
Business & Economics	0.61	1.62	3.15	2.30	1.40	0.00	2.91	3.01	3.15
Arts & Humanities	0.00	1.56	1.07	2.73	1.32	1.85	2.73	2.89	2.92
Other Degrees	0.00	0.40	1.66	2.84	2.47	1.35	2.84	2.88	3.01
All Degrees	0.91	1.37	1.73	1.96	1.59	1.31	2.87	2.96	3.04

Note: Based on a sample of 100,000 simulated observations.

Table 3: Median and Mean Skill Levels

I now turn to a more detailed analysis of skill inequality, looking at the dispersion of skills around subject-specific means or medians. Summary statistics of skill inequality are presented in Table 4. Focussing on the Gini, we can see that mathematical skills tended to be more unequally distributed than verbal skills in the mid 1990s to early 2000s, but this trend has seen a reversal over the time period under consideration. Mathematical skill inequality fell from 53 gini points to 45 and then dropped further to 38, by the end of the sampled period. Conversely, verbal skill inequality, which starts off at a comparatively low level comparable to mathematical skill rises from 24 to 37 and later to 46 gini points, suggesting a large increase of skill inequality across this dimension.

The general trend over the time period seems to go in the direction of greater equality, with overall skill inequality falling from 6 gini points to 5 gini points by the mid-2010s. This result suggests that increases in verbal skill inequality have been more than compensated for by a more equitable distribution of mathematical skills.

These findings are echoed when we look at the square of the coefficient of variation, but we can glean additional insights using the well know decomposition approach. Decomposing the coefficient of variation into between- and within-subject components, I find - unsurprisingly - that the majority of the variation of both skills comes from differences between subjects. Interestingly, however, while the split between the within and between components remains roughly constant over time for mathematical skills, the within share of verbal skills is falling from around 11% to just under 2%, with a corresponding increase in the share of between-subject variation.

The third part of Table 4 summarizes the correlation coefficients for technical skills with organisatorial skills and of both types of skills with overall skills. As per construction within subjects both skills are uncorrelated. At the overall level, however, there is a large negative correlation between both skill types. The overall correlation starts at around

-0.93 in the first time period and then becomes gradually more negative, reaching -0.97 in the final period. Together with the decomposition results, the picture that emerges is one of increasing specialisation: Both STEM and Business & Economics have become almost exclusively focussed on mathematical/technical skills, while Medical and Life Sciences have gone the other way. At the same time, Arts & Humanities and Other Degrees have taken on a more "generalist" position with increasing levels of mathematical/technical skills, but this was not enough to offset the overall trend.

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Gini coefficient	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.098	1.000	0.134	0.072	0.050	0.048	0.061	0.050	0.045
STEM	0.075	0.054	0.046	0.111	0.402	0.001	0.063	0.054	0.046
Business & Economics	0.126	0.096	0.055	0.073	0.021	0.001	0.064	0.053	0.055
Arts & Humanities	0.027	0.070	0.127	0.068	0.079	0.043	0.068	0.052	0.054
Other Degrees	0.003	0.128	0.085	0.059	0.057	0.075	0.059	0.053	0.058
All Degrees	0.531	0.451	0.379	0.243	0.370	0.456	0.065	0.054	0.054

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Squared coefficient of variation	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.031	10514.468	0.059	0.017	0.008	0.007	0.012	0.008	0.006
STEM	0.018	0.009	0.007	0.040	0.750	0.000	0.013	0.009	0.007
Business & Economics	0.051	0.030	0.009	0.017	0.001	0.000	0.013	0.009	0.009
Arts & Humanities	0.002	0.015	0.052	0.015	0.020	0.006	0.015	0.009	0.009
Other Degrees	0.000	0.053	0.023	0.011	0.010	0.018	0.011	0.009	0.011
All Degrees	0.906	0.639	0.443	0.180	0.432	0.658	0.013	0.009	0.009
Between	95.42%	96.44%	95.71%	89.22%	96.82%	98.17%	4.97%	6.08%	10.32%
Within	4.58%	3.56%	4.29%	10.78%	3.18%	1.83%	95.03%	93.92%	89.68%

	Corr(Math, Verbal)			Corr(Math, All)			Corr(Verbal, All)		
Correlation coefficient	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	-0.006	-0.006	0.005	0.808	-0.006	0.325	0.585	1	0.947
STEM	0.012	-0.01	0.001	0.844	0.997	1	0.546	0.071	0.001
Business & Economics	0.001	-0.005	0.002	0.422	0.984	1	0.907	0.176	0.002
Arts & Humanities	-0.008	0.003	-0.006	-0.008	0.722	0.866	1	0.694	0.495
Other Degrees	0	0.016	0	0	0.361	0.813	1	0.938	0.582
All Degrees	-0.925	-0.966	-0.968	0.297	0.304	0.423	0.087	-0.047	-0.184

Note: Based on a sample of 100,000 simulated observations. Between and within variance decomposition based on the square of the coefficient of variation.

Table 4: Skill Inequality

6.3 Wage decomposition

Changes in the distribution of skills and the wider structure of the economy necessarily lead to changes in how productivity accrues to different parts of a worker's human capital. To assess changes over time, I decompose graduate's wages into their constituent parts, using the log wage equation (3.5).

Table 5 presents the share of a worker's log wage that is on average due to technical skills, organisatorial skills and any residual factors, such as their chosen occupation or work experience. Changes are noticeable with respect to the share due to the different types of skills: There is an evident trend with technical skills gaining at the expense of organisatorial skills reflecting the changes in the underlying skill endowments as well as changing skill prices. These effects are sizeable on aggregate, with organisatorial skills accounting for around half of a graduate's wage in the first time period falling to around a third by the mid-2010s. Conversely, technical skills gained in importance, increasing their share from around a quarter to 50%.

I further decompose the variance of the log wage, where a similar trend emerges: across most subjects, with the exception of Medical and Life Sciences, the contribution of technical skills to the variance of log wages has been increasing over time, becoming the dominant factor driving wage dispersion at the subject level. Since, as we saw in the preceding subsection, technical skill inequality has been decreasing, this suggests itself as a driver of the decreasing wage inequality among our sample that we encountered in Figure 1.

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			Residual		
Mean decomposition	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	43%	0%	9%	45%	91%	82%	12%	9%	9%
STEM	60%	90%	91%	27%	1%	0%	13%	9%	9%
Business & Economics	18%	48%	91%	70%	43%	0%	12%	9%	9%
Arts & Humanities	0%	48%	32%	87%	43%	59%	13%	9%	9%
Other Degrees	0%	12%	49%	88%	79%	42%	12%	9%	9%
All Degrees	27%	41%	50%	61%	50%	41%	13%	9%	9%

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			Residual		
Variance decomposition	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	40%	0%	4%	29%	65%	89%	32%	35%	7%
STEM	52%	74%	117%	18%	0%	0%	30%	26%	-17%
Business & Economics	10%	47%	124%	73%	9%	0%	17%	44%	-24%
Arts & Humanities	0%	28%	38%	94%	28%	39%	6%	44%	23%
Other Degrees	0%	4%	52%	87%	54%	32%	13%	41%	16%
All Degrees	314%	602%	683%	330%	603%	635%	-544%	-1105%	-1219%

	Corr(Maths, Wage)			Corr(Verbal, Wage)			Corr(All, Wage)		
Correlation coefficient	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.541	-0.006	0.184	0.391	0.626	0.571	0.669	0.626	0.599
STEM	0.564	0.659	0.614	0.381	0.037	0.004	0.677	0.661	0.614
Business & Economics	0.279	0.641	0.669	0.625	0.121	-0.005	0.684	0.653	0.669
Arts & Humanities	-0.001	0.446	0.572	0.701	0.455	0.322	0.701	0.636	0.658
Other Degrees	-0.002	0.214	0.547	0.66	0.596	0.413	0.66	0.63	0.685
All Degrees	0.141	0.101	0.178	0.122	0.068	-0.017	0.679	0.64	0.643

Note: Based on a sample of 100,000 simulated observations. Mean and variance decompositions are based on the logarithm of wages.

Table 5: Wage Decomposition

7 Counterfactuals

In this section, I consider four counterfactual experiments, in order to assess the importance of different parts of the model for the changing labour market outcomes of graduates. First, I consider the role of within-subject skill inequality. Secondly, I consider a counterfactual world, in which the subject-specific skill distributions do not change across time. Thirdly, I turn the second counterfactual on its head by only allowing the skill distributions to vary while keeping the rest of the economic structure fixed. And finally, I assess the role that changing demographic composition of graduates has played in affecting labour market outcomes. The results from these counterfactuals are summarized in Table 6.

7.1 The role of within-subject skill inequality

This model has been built to incorporate skill differences conditional on degree subject in order to better capture the considerable degree of within-subject heterogeneity presented at the beginning of this paper. In this subsection, I consider how large a role is played by allowing for heterogenous skills within subject categories. Ex-ante, one might expect two effects: additional skill heterogeneity will **i)** increase inequality of wage outcomes, and (perhaps less obviously) **ii)** increase mean wages in an environment where skills are multidimensional and graduates are able to choose their utility maximizing occupation.

To assess the role of within-subject skill inequality, I set $\sigma^2_{kmt} = 0, \forall k, m, t$ and then re-simulate the model. This is equivalent to assuming that each graduate has the median amount of each skill within his degree subject cohort. Table 6 summarizes the impact of this adjustment on mean hourly wages and wage inequality.[22]

[22] Occupational outcomes are barely affected by the removal of within-subject heterogeneity, which is likely due to the specific model setup. Results are presented nonetheless for completeness.

Counterfactual I - No within degree skill heterogeneity

	Mean Wage (Δ %)			Gini Wage (Δ %)			HHI (Δ %)			Nontypical Share (Δ %)		
Subject	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019
Medical and Life Sciences	-4.69	-2.66	-2.95	-27.71	-22.95	-21.54	-0.15	0.04	-0.17	0.29	-0.29	0.33
STEM	-5.29	-3.83	-2.56	-28.72	-26.44	-22.25	-0.41	-0.22	-0.06	1.03	0.44	0.12
Business & Economics	-4.97	-3.71	-3.26	-29.40	-25.05	-27.63	-0.15	-0.19	-0.33	0.40	0.52	0.66
Arts & Humanities	-4.64	-3.88	-3.88	-31.04	-24.30	-26.20	-0.29	0.00	-0.13	0.81	-0.23	0.52
Other Degrees	-3.50	-3.71	-4.57	-26.34	-23.48	-29.55	-0.05	-0.22	-0.21	0.25	0.66	0.28
All Degrees	-4.70	-3.53	-3.45	-28.44	-23.89	-24.64	-0.23	-0.11	-0.17	0.61	0.24	0.33

Counterfactual II - Skill distribution parameters fixed

	Mean Wage (Δ %)			Gini Wage (Δ %)			HHI (Δ %)			Nontypical Share (Δ %)		
Subject	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019
Medical and Life Sciences	0.00	-6.57	-8.64	0.00	5.41	10.62	0.00	-1.38	-0.66	0.00	1.45	0.59
STEM	0.00	-4.52	-10.39	0.00	5.48	8.56	0.00	0.44	-0.44	0.00	-0.19	2.48
Business & Economics	0.00	-3.12	-7.06	0.00	10.41	2.74	0.00	0.20	-0.48	0.00	3.53	8.86
Arts & Humanities	0.00	-4.38	-8.76	0.00	13.10	7.71	0.00	0.59	-0.50	0.00	3.71	4.74
Other Degrees	0.00	-0.74	-5.98	0.00	4.24	-2.66	0.00	0.20	-0.29	0.00	-0.80	-0.96
All Degrees	0.00	-4.12	-8.11	0.00	7.75	4.80	0.00	0.00	-0.47	0.00	1.44	3.51

Counterfactual III - Other labour market factors fixed

	Mean Wage (Δ %)			Gini Wage (Δ %)			HHI (Δ %)			Nontypical Share (Δ %)		
Subject	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019
Medical and Life Sciences	0.00	-9.61	-5.33	0.00	9.36	3.05	0.00	4.45	2.30	0.00	-11.87	-2.12
STEM	0.00	-8.33	-2.59	0.00	11.29	3.23	0.00	5.41	1.51	0.00	-10.55	1.12
Business & Economics	0.00	-8.78	-2.50	0.00	10.67	2.71	0.00	5.14	1.09	0.00	-17.69	-6.93
Arts & Humanities	0.00	-8.66	-4.44	0.00	11.05	2.87	0.00	4.86	1.75	0.00	-16.69	-7.17
Other Degrees	0.00	-9.15	-3.97	0.00	9.75	2.20	0.00	4.86	1.51	0.00	-11.09	3.81
All Degrees	0.00	-8.90	-3.89	0.00	10.29	3.09	0.00	4.96	1.69	0.00	-16.98	-7.37

Counterfactual IV - Demographic composition fixed

	Mean Wage (Δ %)			Gini Wage (Δ %)			HHI (Δ %)			Nontypical Share (Δ %)		
Subject	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019	1994 - 2001	2002 - 2011	2012 - 2019
Medical and Life Sciences	0.00	0.36	0.72	0.00	-0.73	-1.21	0.00	0.63	2.06	0.00	-3.64	-3.45
STEM	0.00	0.12	0.84	0.00	0.08	-0.29	0.00	0.40	0.13	0.00	-0.35	-1.33
Business & Economics	0.00	-0.17	0.88	0.00	0.41	1.08	0.00	-1.01	-2.35	0.00	4.74	4.72
Arts & Humanities	0.00	1.38	2.52	0.00	0.08	0.02	0.00	0.23	-0.13	0.00	-1.07	5.86
Other Degrees	0.00	0.36	0.48	0.00	0.91	-1.11	0.00	0.03	-0.64	0.00	0.28	-2.15
All Degrees	0.00	0.47	1.45	0.00	0.01	-0.03	0.00	-0.07	-0.04	0.00	-0.72	-1.09

Note: Percentage deviations are calculated relative to baseline estimates.

Table 6: Model Counterfactuals

The results suggest that within-subject skill inequality plays a large role in explaining graduate wage inequality within-subject and overall. The gini coefficient for hourly wages from 1994 to 2002 is around 28% lower compared to the data. In the period 2003 to 2019 this drops to around 24% which is lower but still constitutes a sizeable effect. For mean wages, the effects are less pronounced than expected, but still imply a reduction of 4.7% in mean hourly wages in the first period. Later on this impact drops to around 3.5% which is still sizeable. As we saw in the preceding section, overall skill inequality has fallen across these time periods thereby reducing the impact of within-subject skill heterogeneity on labour market outcomes. However, within-subject skill heterogeneity does still play a sizeable role in wage dynamics even in the most recent sample.

7.2 The role of a changing skill distribution

In this next counterfactual, I assess the impact of the changing skill distribution on the outcomes of graduates. For this purpose, I fix both μ_{kmt} and σ^2_{kmt} at their period 1 values while allowing all other parameters to change across periods. Table 6 summarizes the effect of this adjustment on the labour market outcomes of graduates:

Fixing the subject-specific skill distributions implies that overall changes in the labour market outcomes of graduates are due to **i)** changes in the composition of graduates, and **ii)** changes in the other structural parts of the labour market. Considering periods 2003 - 2011 and 2012 - 2019, the counterfactual analysis suggests that mean wages would have fallen considerably across all subjects, with an average decrease of 4.12% in the second period and a drop of 8.11% in the third.

The effect is of similar size when considering wage inequality. Here the counterfactual predicts a gini coefficient that is about 7.75% higher than that observed in the data in 2003 - 2011 and 4.8% higher in 2012 -2019, suggesting that the overall decline in skill inequality that I discussed in the previous section had a considerable impact in reducing graduate level wage inequality.

The impact of the changing skill distribution does not seem to have a large effect on occupation choices. The HH index of occupational concentration is almost unaffected compared to the baseline model.[23] Although the share of non-typical occupations is increased by 1.44% in 2003 - 2011 and 3.51% in 2012 - 2019 than predicted by the baseline model, suggesting that this mechanism has a role to play here also.

[23] I suggest that this is because the changes in the within-subject skill distributions are not large enough to effect much of a change in preferred occupations. Instead, the occupation distribution appears to be driven by changes in the composition of graduates and other external factors.

7.3 The role of changing task weights

Complementing the preceding analysis I perform another counterfactual simulation, this time fixing the occupation-specific skill weights λ_o at their $t = 1$ values while letting all other parameters change across periods. This provides an assessment of the importance of changes in the evolving skill requirements of firms. The results are again presented in Table 6.

The results suggest that had firms continued to operate with their initial technologies, average wages would be around 8.9% lower in 2003 - 2011 and 3.89% lower in 2012 - 2019. When graduate's skills are not well matched to the demands of employers wages decline, supporting the hypothesis that part of the increase in mathematical skills that I observe was driven in response to changing demands of firms.[24] Wage inequality would also have been larger in the absence of changes to the occupation task weights, with the wage gini being increased by up to 10% relative to the baseline model.

Interestingly, the changing skill demands appear to have played a larger role for occupational choices than the changing skill distributions I considered before. Particularly the share of graduates entering non-typical occupations decreases by 16.98% in 2003 - 2011 and by 7.37% in 2012 - 2019, suggesting that this channel is important for explaining this particular trend.

7.4 The role of changing demographics

Finally, I assess the role that changing demographics played in graduates' labour market outcomes. For this purpose, I fix the distribution of demographic characteristics (experience, sex, subject choice) to 1994-2002 while allowing all parameters to vary over time. The results of this counterfactual analysis are presented in the last section of Table 6. In line with the observation, that demographic characteristics did not change significantly over the period[25], there are only minor changes to the graduates' outcomes. The one notable finding is that mean wages would be around 0.5% higher in 2003-2011 and almost 1.5% higher in 2012-2019. This is probably reflective of a higher share of high-paid STEM and Business & Economics graduates in the initial period.

[24] The alternative hypothesis that employers changed their production technology in response to changing skill supplies, cannot be ruled out but appears to be less plausible.

[25] See Table A3 for details.

8 Conclusion

The formation of human capital and the acquisition of specific skills lies at the heart of a university education. Despite their centrality to both the academic and public discourse about tertiary education, there are few quantitative studies that actually investigate what skills graduates possess at the end of their university courses. This paper has been an attempt at estimating the UK's graduate skill distribution and changes to it over the last 3 decades.

My findings suggest, that while there has been a considerable change in the composition of skills that university graduates possess, this has not resulted in a large decline in the overall skill level of a typical graduate. While verbal/organisatorial skills have decreased, the typical graduate in the most recent time period has a lot more technical skills than previous cohorts. The observed changes are due to both changes in the composition of graduates and changes in the skill distribution at the subject level. Overall skill inequality has declined in the wake of the expansion of higher education, making graduates more homogenous in terms of skillsets over time. I speculate that this pattern is a result of the increase in the demand for mathematical skills in the wake of increasing use of ICT in recent decades, as well as the UK government's efforts to encourage technical skills formation throughout the early 2000s.

Taking together, mathematical skills now contribute around 50% to the overall monetary compensation of the average graduate. The combination of increasing importance of mathematical/technical skills and their more equal distribution has led to a fall in residual wage inequality amongst recent cohorts of graduates in their early careers.

Using counterfactual simulations, I find that in the absence of changes in the subject-specific skill distributions, real hourly wages would be up to 8% lower in 2012 - 2019. Additionally, wage inequality would also be larger by around 5%, compared to what is observed in the data. Further analysis suggests, that the changing demographic composition of graduates has played only a minor when it comes to changing labour market outcomes.

Overall the results of the paper suggest that university graduates enter the labour market with more technical and mathematical abilities, which helps them perform in an economy that is increasingly relying on these skills. On one level this speaks to the success of the UK's education sector in producing graduates with highly relevant skills and abilities, despite the large and ongoing challenges of the last decades. On a more cautionary point, increasing specialisation and reliance on technical skills at the expense of "softer skills", might spell issues for recent cohorts of graduates as they progress in their careers. I leave this and other questions for future research.

BIBLIOGRAPHY

[1] Acemoglu, Daron, and David Autor. "Skills, tasks and technologies: Implications for employment and earnings." In Handbook of labor economics, vol. 4, pp. 1043-1171. Elsevier, 2011.

[2] Altonji, Joseph G., Peter Arcidiacono, and Arnaud Maurel. "The analysis of field choice in college and graduate school: Determinants and wage effects." In Handbook of the Economics of Education, vol. 5, pp. 305-396. Elsevier, 2016.

[3] Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer. "Trends in earnings differentials across college majors and the changing task composition of jobs." American Economic Review 104, no. 5 (2014): 387-393.

[4] Andrews, Rodney J., Scott A. Imberman, Michael F. Lovenheim, and Kevin M. Stange. The Returns to College Major Choice: Average and Distributional Effects, Career Trajectories, and Earnings Variability. No. w30331. National Bureau of Economic Research, 2022.

[5] Autor, David, H. "The “task approach” to labor markets: an overview." Journal for Labour Market Research 46, no. 3 (2013): 185-199.

[6] Autor, David H., and Michael J. Handel. "Putting tasks to the test: Human capital, job tasks, and wages." Journal of labor Economics 31, no. S1 (2013): S59-S96.

[7] Autor, David, and F. Levy. "R. Murnane, 2003, The Skill Content of Recent Technological Change: An Empirical Exploration." Quarterly Journal of Economics 118, no. 4.

[8] Bisello, Martina. "Job polarization in Britain from a task-based perspective. Evidence from the UK Skills Surveys." Department of Economics and Management, University of Pisa Discussion Paper 160 (2013).

[9] Blundell, Richard, David Green, and Wenchao Jin. "The UK education expansion and technological change." Institute for Fiscal Studies Working Paper (2018).

[10] Gill, Jeff, and Gary King. "What to do when your Hessian is not invertible: Alternatives to model respecification in nonlinear estimation." Sociological methods & research 33, no. 1 (2004): 54-87.

[11] Goldin, Claudia, and Lawrence F. Katz. The race between education and technology. harvard university press, 2010.

[12] Gourieroux, Christian, and Alain Monfort. "Simulation-based inference: A survey with special reference to panel data models." Journal of Econometrics 59, no. 1-2 (1993): 5-33.

[13] Hajivassiliou, Vassilis A., and Paul A. Ruud. "Classical estimation methods for LDV models using simulation." Handbook of econometrics 4 (1994): 2383-2441.

[14] Hamermesh, Daniel S., and Stephen G. Donald. "The effect of college curriculum on earnings: An affinity identifier for non-ignorable non-response bias." Journal of Econometrics 144, no. 2 (2008): 479-491.

[15] Hastings, Justine S., Christopher A. Neilson, and Seth D. Zimmerman. Are some degrees worth more than others? Evidence from college admission cutoffs in Chile. No. w19241. National Bureau of Economic Research, 2013.

[16] Hemelt, Steven W., Brad Hershbein, Shawn M. Martin, and Kevin M. Stange. College majors and skills: evidence from the universe of online job ads. No. w29605. National Bureau of Economic Research, 2021.

[17] Holmes, Craig, and Ken Mayhew. "The economics of higher education." Oxford review of economic policy 32, no. 4 (2016): 475-496.

[18] Kinsler, Josh, and Ronni Pavan. "The specificity of general human capital: Evidence from college major choice." Journal of Labor Economics 33, no. 4 (2015): 933-972.

[19] Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad. "Field of study, earnings, and self-selection." The Quarterly Journal of Economics 131, no. 3 (2016): 1057-1111.

[20] Lee, Lung-Fei. "Asymptotic bias in simulated maximum likelihood estimation of discrete choice models." Econometric Theory 11, no. 3 (1995): 437-483.

[21] Leighton, Margaret, and Jamin D. Speer. "Labor market returns to college major specificity." European Economic Review 128 (2020): 103489.

[22] Lindley, Joanne, and Steven McIntosh. "Growth in within graduate wage inequality: The role of subjects, cognitive skill dispersion and occupational concentration." Labour Economics 37 (2015): 101-111.

[23] Lovenheim, Michael F., and Jonathan Smith. Returns to different postsecondary investments: Institution type, academic programs, and credentials. No. w29933. National Bureau of Economic Research, 2022.

[24] McFadden, Daniel. "The measurement of urban travel demand." Journal of public economics 3, no. 4 (1974): 303-328.

[25] McFadden, Daniel, and Kenneth Train. "Mixed MNL models for discrete response." Journal of applied Econometrics 15, no. 5 (2000): 447-470.

[26] Onozuka, Yuki. "Heterogeneous Skill Growth across College Majors." (2019): 1-58.

[27] Rohrbach-Schmidt, Daniela, and Michael Tiemann. "Changes in workplace tasks in Germany—evaluating skill and task measures." Journal for Labour Market Research 46, no. 3 (2013): 215-237.

[28] Roy, Andrew Donald. "Some thoughts on the distribution of earnings." Oxford economic papers 3, no. 2 (1951): 135-146.

[29] Roys, Nicolas, and Christopher Taber. "Skills prices, occupations and changes in the wage structure." (2016).

[30] Sanders, Carl, and Christopher Taber. "Life-cycle wage growth and heterogeneous human capital." Annu. Rev. Econ. 4, no. 1 (2012): 399-425.

[31] Train, Kenneth. "Halton sequences for mixed logit." (2000).

[32] Train, Kenneth E. "EM algorithms for nonparametric estimation of mixing distributions." Journal of Choice Modelling 1, no. 1 (2008): 40-69.

[33] Train, Kenneth E. Discrete choice methods with simulation. Cambridge university press, 2009.

9 Appendix A

9.1 Model estimates

The following tables presents the SMLE results.

	Medical & Life Sciences	STEM	Business & Economics	Arts & Humanities	Other Degrees
			1994 - 2002		
μ_{Math}	0.363	0.711	-0.51	-11.409	-11.433
	(1.434)	(1.365)	(3.325)	(480.107)	(745.178)
μ_{Verbal}	0.359	-0.14	0.825	0.998	1.038
	(1.422)	(2.067)	(1.683)	(0.953)	(1.441)
σ_{Math}	0.172	0.133	0.223	0.048	0.005
	(0.47)	(0.348)	(1.159)	(191.509)	(148.002)
σ_{Verbal}	0.128	0.197	0.131	0.12	0.103
	(0.312)	(0.639)	(0.493)	(0.325)	(0.457)
			2003 - 2011		
μ_{Math}	-85.494	1.102	0.468	0.439	-0.934
	(0)	(1.426)	(1.013)	(1.116)	(2.22)
μ_{Verbal}	1.074	-3.877	0.334	0.269	0.9
	(0.191)	(16.006)	(1.078)	(1.214)	(0.888)
σ_{Math}	8.277	0.095	0.17	0.123	0.226
	(0)	(0.314)	(0.374)	(0.296)	(0.72)
σ_{Verbal}	0.089	0.746	0.036	0.139	0.102
	(0.071)	(7.171)	(0.365)	(0.357)	(0.22)
			2012 - 2019		
μ_{Math}	-1.181	1.148	1.144	0.047	0.498
	(1.751)	(0.472)	(0.616)	(0.757)	(0.525)
μ_{Verbal}	0.977	-11.606	-8.449	0.613	0.29
	(0.581)	(271.308)	(72.858)	(0.553)	(0.572)
σ_{Math}	0.238	0.082	0.097	0.224	0.149
	(0.592)	(0.133)	(0.187)	(0.317)	(0.155)
σ_{Verbal}	0.084	0.001	0.001	0.076	0.133
	(0.123)	(90.377)	(97.687)	(0.075)	(0.135)

Numerical Standard Errors in Parentheses

Table A1: Simulated Maximum Likelihood Estimates - Skill Distribution Parameters

9.2 Additional Tables

	1994 - 2002	2003 - 2011	2012 - 2019
ε	0.076	0.057	0.045
	(0.041)	(0.033)	(0.028)
γ	-0.019	0.000	-0.043
	(0.045)	(0.031)	(0.028)
η_1	0.000	0.000	0.000
	-	-	-
η_2	0.045	0.146	0.072
	(0.187)	(0.264)	(0.085)
η_3	0.150	0.154	0.149
	(0.259)	(0.301)	(0.094)
η_4	0.116	0.148	0.137
	(0.277)	(0.176)	(0.071)
η_5	0.307	0.095	0.142
	(0.613)	(0.626)	(0.194)
η_6	0.110	0.055	0.068
	(0.517)	(0.431)	(0.238)
η_7	0.333	0.207	0.439
	(0.332)	(0.195)	(0.096)
η_8	0.445	0.236	0.280
	(0.353)	(0.218)	(0.066)
η_9	0.455	0.217	0.595
	(0.239)	(0.251)	(0.149)
ω_1	0.000	0.000	0.000
	-	-	-
ω_2	0.869	1.040	0.702
	(0.122)	(0.149)	(0.144)
ω_3	0.737	1.223	1.252
	(0.162)	(0.121)	(0.137)
ω_4	0.756	0.954	0.614
	(0.104)	(0.134)	(0.144)
ω_5	-1.743	-1.583	-1.515
	(0.287)	(0.388)	(0.195)
ω_6	-0.968	-0.272	-0.194
	(0.25)	(0.199)	(0.137)
ω_7	-0.377	0.395	0.055
	(0.142)	(0.156)	(0.161)
ω_8	-1.841	-1.867	-1.634
	(0.322)	(0.192)	(0.176)
ω_9	-1.203	-0.310	-0.286
	(0.292)	(0.139)	(0.129)
τ_1	0.000	0.000	0.000
	-	-	-
τ_2	-0.027	-0.001	-0.021
	(0.09)	(0.136)	(0.034)
τ_3	0.003	-0.006	-0.027
	(0.113)	(0.101)	(0.033)
τ_4	-0.004	0.014	0.025
	(0.112)	(0.086)	(0.039)
τ_5	0.068	0.065	0.033
	(0.112)	(0.05)	(0.044)
τ_6	0.088	0.032	0.038
	(0.121)	(0.068)	(0.044)
τ_7	0.126	0.030	0.040
	(0.125)	(0.063)	(0.032)
τ_8	0.159	-0.018	0.052
	(0.163)	(0.046)	(0.033)
τ_9	0.169	-0.044	-
	(0.114)	(0.074)	-

Numerical Standard Errors in Parentheses

Table A2: Simulated Maximum Likelihood Estimates - Other Parameters

	Share of graduates (%)			Share of female of graduates (%)			Average experience		
	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	16.84	21.94	26.29	67.38	71.66	70.72	1.19	1.18	1.23
STEM	26.06	21.19	19.8	27.34	27.76	29.62	1.25	1.21	1.28
Business & Economics	20.26	19.49	13.87	50.38	56.22	55.26	1.25	1.28	1.2
Arts & Humanities	19.43	22.33	14.21	61.33	64.17	57.42	1.33	1.21	1.19
Other Degrees	17.41	15.05	25.82	67.71	66.85	62.62	1.25	1.22	1.28
All Degrees	3860	3587	3222	52.38	56.96	56.46	1.25	1.22	1.24

Table A3: Summary statistics of QLFS sample - Demographics

	Mean hourly wage			Gini hourly wage			HHI index of occupation concentration			Share of nontypical occupations		
	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	10.61	11.54	10.45	0.17	0.17	0.16	0.29	0.25	0.25	12.15	18.93	20.43
STEM	10.81	11.6	11.4	0.18	0.17	0.16	0.25	0.26	0.26	12.72	16.45	15.52
Business & Economics	10.18	10.87	10.41	0.18	0.17	0.17	0.22	0.22	0.21	4.22	7.44	7.38
Arts & Humanities	9.2	9.66	9.3	0.19	0.17	0.16	0.19	0.19	0.19	10.4	14.36	16.38
Other Degrees	10.28	10.86	9.86	0.18	0.18	0.18	0.25	0.25	0.17	9.38	11.48	20.79
All Degrees	10.24	10.9	10.32	0.18	0.18	0.17	0.21	0.2	0.2	8.42	10.93	13.78

Note: Wages are CPI deflated (2014 = 100).

Table A4: Summary statistics of QLFS sample - Labour Market Outcomes

Mathematical/Technical Tasks

Variable name	Description
cspecial	importance of: specialist knowledge or understanding
cfaults	importance of: spotting problems or faults
csolutn	importance of: thinking of solutions to problems
canalyse	importance of: analysing complex problems in depth
ccalca	importance of: arithmetic (adding, subtracting, multiplying, dividing)
cstats	importance of: advanced mathematics/ statistics
cpercent	importance of: arithmetic involving fractions (decimals, percentages, fractions)

Verbal/Organisatorial Tasks

Variable name	Description
cteach	importance of: teaching people (individuals or groups)
cspeech	importance of: making speeches/ presentations
cteamwk	importance of: working with a team
corgwork	importance of: knowledge of how organisation works
cplanme	importance of: planning own activities
cplanoth	importance of: planning the activities of others
cmytime	importance of: organising own time
cahead	importance of: thinking ahead
cread	importance of: reading written information (eg. forms, notices, signs)
cshort	importance of: reading short documents
clong	importance of: reading long documents
cwrite	importance of: writing materials such as forms, notices or signs
cwritesh	importance of: writing short documents
cwritelg	importance of: writing long documents

Table A5: Variables used in the construction of task intensity measures

	Mathematical-Technical Tasks			Verbal-Organisatorial Tasks		
	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Managers and senior officials	0.73	0.75	0.73	0.75	0.77	0.76
Professional occupations	0.75	0.76	0.75	0.78	0.79	0.80
Associate professional and technical occupations	0.68	0.71	0.68	0.71	0.76	0.73
Administrative and secretarial occupations	0.61	0.67	0.65	0.64	0.69	0.69
Skilled trades occupations	0.64	0.71	0.69	0.56	0.63	0.63
Personal service occupations	0.53	0.60	0.58	0.62	0.72	0.69
Sales and customer service occupations	0.56	0.62	0.53	0.54	0.60	0.53
Process, plant and machine operatives	0.52	0.61	0.57	0.49	0.60	0.58
Elementary occupations	0.42	0.56	0.42	0.44	0.59	0.47
Mean	0.67	0.70	0.67	0.70	0.73	0.70
Standard Deviation	0.08	0.06	0.09	0.08	0.06	0.09

Table A6: Task Intensities

9.3 Additional Figures

Figure A1: Visualization of the graduate skill distribution

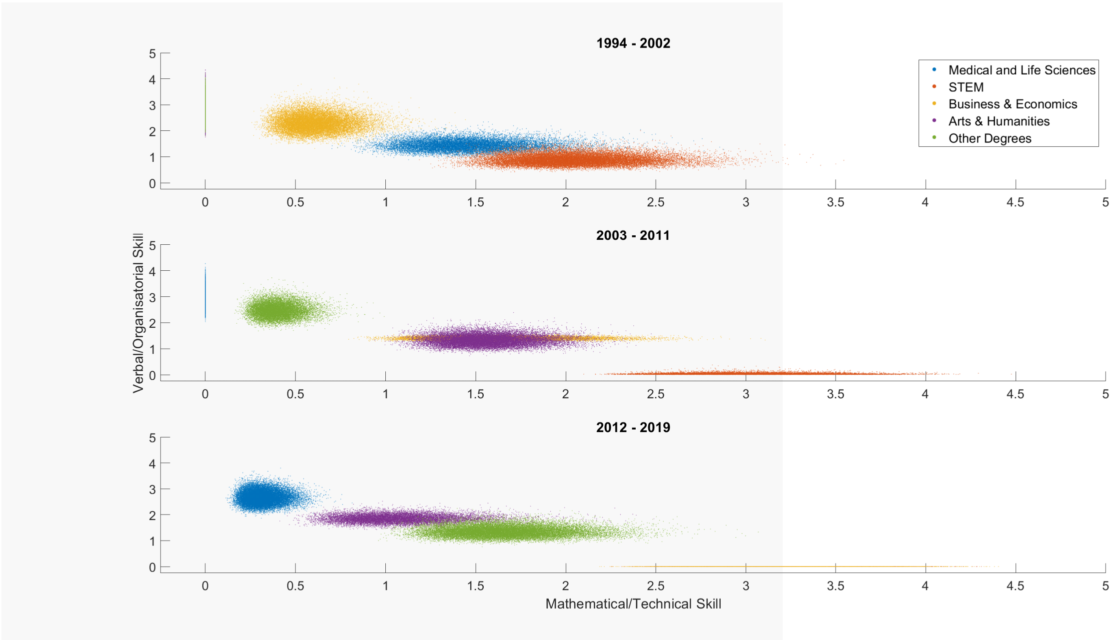

Note: Based on a simulated sample of 100,000 observations.

10 Appendix B

10.1 Postgraduate qualifications

This subsection summarizes the key outcomes for the model, when postgraduate qualifications are included in the sample.

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Median Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.391	0.004	0.697	1.479	2.963	2.315	2.886	2.967	3.022
STEM	1.880	2.627	3.197	1.041	0.428	0.000	2.938	3.066	3.197
Business & Economics	0.003	1.792	3.149	2.876	1.232	0.000	2.879	3.037	3.149
Arts & Humanities	0.003	1.397	0.949	2.722	1.481	1.955	2.725	2.891	2.915
Other Degrees	0.007	0.002	1.497	2.834	2.883	1.505	2.841	2.885	3.018
All Degrees	0.007	1.407	1.474	2.433	1.468	1.536	2.861	2.971	3.068

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Mean Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.412	0.004	0.708	1.492	2.975	2.324	2.904	2.978	3.032
STEM	1.898	2.641	3.209	1.058	0.438	0.000	2.956	3.079	3.209
Business & Economics	0.003	1.813	3.166	2.896	1.239	0.000	2.898	3.052	3.166
Arts & Humanities	0.003	1.409	0.976	2.741	1.493	1.961	2.744	2.902	2.937
Other Degrees	0.007	0.002	1.518	2.850	2.896	1.520	2.857	2.898	3.038
All Degrees	0.757	1.222	1.843	2.121	1.763	1.236	2.878	2.985	3.079

Note: Based on a sample of 100,000 simulated observations.

Table B1: Median and Mean Skill Levels - Estimation including Postgraduate Qualifications

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Gini coefficient	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.097	0.000	0.092	0.074	0.050	0.052	0.061	0.050	0.046
STEM	0.077	0.056	0.045	0.100	0.124	0.001	0.062	0.052	0.045
Business & Economics	0.001	0.080	0.057	0.066	0.055	0.001	0.066	0.053	0.057
Arts & Humanities	0.001	0.074	0.137	0.065	0.073	0.045	0.065	0.052	0.055
Other Degrees	0.000	0.000	0.093	0.059	0.051	0.080	0.059	0.051	0.061
All Degrees	0.607	0.489	0.328	0.228	0.329	0.440	0.064	0.053	0.055

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Squared CV	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.030	0.000	0.027	0.017	0.008	0.009	0.012	0.008	0.007
STEM	0.019	0.010	0.006	0.032	0.049	0.000	0.012	0.008	0.006
Business & Economics	0.000	0.021	0.010	0.014	0.009	0.000	0.014	0.009	0.010
Arts & Humanities	0.000	0.017	0.062	0.014	0.017	0.006	0.014	0.009	0.010
Other Degrees	0.000	0.000	0.027	0.011	0.008	0.020	0.011	0.008	0.012
All Degrees	1.346	0.768	0.346	0.162	0.337	0.639	0.013	0.009	0.010
Between	96.38%	96.89%	94.95%	88.86%	96.07%	97.62%	4.78%	6.99%	10.20%
Within	3.62%	3.11%	5.05%	11.14%	3.93%	2.38%	95.22%	93.01%	89.80%

Note: Based on a sample of 100,000 simulated observations. Between and within variance decomposition based on the square of the coefficient of variation.

Table B2: Skill Inequality - Estimation including Postgraduate Qualifications

10.2 Changing ϕ

This subsection summarizes the key outcomes for the model, when ϕ is increased or decreased by 10%.

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Median Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.383	0.000	0.254	1.485	2.926	2.716	2.883	2.926	2.979
STEM	1.966	3.037	3.155	0.938	0.000	0.000	2.922	3.037	3.155
Business & Economics	0.402	1.772	3.140	2.453	1.222	0.000	2.880	3.002	3.140
Arts & Humanities	0.000	1.602	0.984	2.714	1.262	1.908	2.714	2.875	2.899
Other Degrees	0.000	0.416	1.605	2.825	2.440	1.382	2.825	2.864	2.997
All Degrees	0.478	1.618	1.538	2.299	1.350	1.434	2.851	2.943	3.030

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Mean Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.402	0.000	0.262	1.494	2.935	2.722	2.896	2.935	2.984
STEM	1.984	3.052	3.164	0.953	0.000	0.000	2.937	3.052	3.164
Business & Economics	0.437	1.775	3.155	2.466	1.245	0.000	2.903	3.019	3.155
Arts & Humanities	0.000	1.610	1.010	2.733	1.277	1.911	2.733	2.887	2.921
Other Degrees	0.000	0.425	1.621	2.839	2.450	1.392	2.839	2.875	3.013
All Degrees	0.848	1.419	1.698	2.019	1.538	1.344	2.867	2.957	3.042

Note: Based on a sample of 100,000 simulated observations.

Table B3: Median and Mean Skill Levels - Higher ϕ value

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Gini coefficient	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.095	0.831	0.130	0.064	0.045	0.042	0.056	0.045	0.040
STEM	0.071	0.049	0.042	0.099	1.000	0.001	0.058	0.049	0.042
Business & Economics	0.227	0.032	0.050	0.057	0.104	0.001	0.060	0.047	0.050
Arts & Humanities	0.090	0.054	0.129	0.063	0.082	0.031	0.063	0.048	0.049
Other Degrees	0.080	0.111	0.083	0.054	0.053	0.064	0.054	0.048	0.054
All Degrees	0.554	0.436	0.393	0.228	0.391	0.453	0.060	0.049	0.050

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Squared coefficient of variation	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.029	31.695	0.055	0.013	0.007	0.005	0.010	0.007	0.005
STEM	0.016	0.008	0.005	0.032	-	0.000	0.011	0.008	0.005
Business & Economics	0.180	0.003	0.008	0.010	0.035	0.000	0.012	0.007	0.008
Arts & Humanities	0.026	0.009	0.054	0.013	0.022	0.003	0.013	0.007	0.008
Other Degrees	0.020	0.040	0.022	0.009	0.009	0.013	0.009	0.007	0.009
All Degrees	1.005	0.603	0.478	0.161	0.476	0.650	0.012	0.008	0.008
Between	95.42%	96.44%	95.71%	89.22%	96.82%	98.17%	4.97%	6.08%	10.32%
Within	4.58%	3.56%	4.29%	10.78%	3.18%	1.83%	95.03%	93.92%	89.68%

Note: Based on a sample of 100,000 simulated observations. Between and within variance decomposition based on the square of the coefficient of variation.

Table B4: Skill Inequality - Higher ϕ values

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Median Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.308	0.000	0.341	1.557	2.924	2.621	2.885	2.924	2.975
STEM	1.884	3.036	3.143	1.016	0.000	0.005	2.923	3.036	3.148
Business & Economics	0.003	1.627	3.131	2.876	1.361	0.000	2.878	2.996	3.131
Arts & Humanities	0.003	1.568	1.090	2.717	1.290	1.797	2.720	2.873	2.902
Other Degrees	0.005	0.417	1.671	2.824	2.433	1.306	2.829	2.862	2.992
All Degrees	0.005	1.480	1.603	2.433	1.403	1.370	2.852	2.940	3.025

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Mean Skills	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	1.332	0.000	0.352	1.571	2.937	2.630	2.902	2.937	2.983
STEM	1.907	3.054	3.157	1.036	0.000	0.005	2.942	3.054	3.162
Business & Economics	0.003	1.651	3.150	2.896	1.365	0.000	2.898	3.016	3.150
Arts & Humanities	0.003	1.583	1.117	2.740	1.305	1.806	2.743	2.888	2.923
Other Degrees	0.005	0.430	1.690	2.844	2.447	1.322	2.848	2.877	3.012
All Degrees	0.729	1.390	1.753	2.143	1.568	1.288	2.872	2.958	3.041

Note: Based on a sample of 100,000 simulated observations.

Table B5: Median and Mean Skill Levels - Lower ϕ value

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Gini coefficient	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.108	0.998	0.140	0.076	0.054	0.053	0.065	0.054	0.050
STEM	0.083	0.057	0.051	0.109	-	0.000	0.067	0.057	0.051
Business & Economics	0.001	0.097	0.058	0.067	0.042	0.001	0.067	0.056	0.058
Arts & Humanities	0.001	0.076	0.125	0.071	0.083	0.053	0.071	0.056	0.058
Other Degrees	0.001	0.135	0.087	0.063	0.061	0.084	0.063	0.056	0.062
All Degrees	0.623	0.449	0.371	0.227	0.379	0.459	0.068	0.058	0.058

	Mathematical/Technical Skill			Verbal/Organisatorial Skill			All Skills		
Squared coefficient of variation	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019	1994 - 2002	2003 - 2011	2012 - 2019
Medical and Life Sciences	0.038	2969.574	0.064	0.019	0.009	0.009	0.013	0.009	0.008
STEM	0.022	0.010	0.008	0.038	-	0.000	0.014	0.010	0.008
Business & Economics	0.000	0.030	0.011	0.014	0.006	0.000	0.014	0.010	0.011
Arts & Humanities	0.000	0.018	0.051	0.016	0.022	0.009	0.016	0.010	0.011
Other Degrees	0.000	0.060	0.024	0.012	0.012	0.023	0.012	0.010	0.012
All Degrees	1.427	0.632	0.422	0.159	0.452	0.663	0.015	0.011	0.011
Between	95.77%	96.05%	95.04%	87.37%	96.33%	97.65%	3.96%	5.34%	8.91%
Within	4.23%	3.95%	4.96%	12.63%	3.67%	2.35%	96.04%	94.66%	91.09%

Note: Based on a sample of 100,000 simulated observations. Between and within variance decomposition based on the square of the coefficient of variation.

Table B6: Skill Inequality - Lower ϕ values

11 Appendix C

11.1 Calibration of ϕ^2

Wave 3 of the Understanding Society Survey contains a module assessing the cognitive and psychological traits of adult (16+) respondents. Questionnaire items include test measuring 1. Numeric Ability, 2. A Subtraction Exercise, 3. Completion of a Number Sequence, 4. A word recall exercise & 5. Verbal Fluency. I use these items to generate 2 measures of skill, mapping into the dimensions of mathematical and verbal skill used in the model, using principal component analysis on the standardized survey responses. Then I use these measures to run a cross-sectional (log) wage regression, controlling for 1 Digit SOC (2000) occupation as well as the full set of interactions with the two skill measures. The resulting regression equation is exactly the proxy analogue to the log wage equation:

$$w_i = \sum_O \mathbf{1}_{(o_i^*=o)} \tilde{\eta}_o + \sum_O \mathbf{1}_{(o_i^*=o)} \tilde{\eta}_o \tilde{s}_{i,math} + \sum_O \mathbf{1}_{(o_i^*=o)} \tilde{\eta}_o \tilde{s}_{i,verbal} + \tilde{\nu}_i$$

The residual variance of $\tilde{\nu}_i$ provides an estimate for ϕ^2.

I estimate the auxiliary model on a sample of full-time working individuals aged 21-27, adding additional controls for sex and age as a proxy of labour market experience. The estimates are presented in the next table:

11.2 Estimation algorithm

The estimation procedure is a simple application of simulated maximum likelihood. In maximum likelihood, we find a vector of parameters so that the model maximizes the probability of observing the actual outcome.

The only complication, that arises here comes from the fact, that we do not have a closed form solution for the joint probability (3.22) and thus have to evaluate the integral via simulation. This can be done by taking draws from the distribution of s, evaluating $\Pr(o_i^*, w_i^{obs}|s_i)$ at each of these draws and then averaging over the results. Standard results suggest, that as long as one uses a large enough number of draws to approximate the integral, the Maximum Simulated Likelihood Estimation (MSLE) is asymptotically equivalent to classical Maximum Likelihood Estimation (MLE) (c.f. McFadden & Train (2000)). For a proof that the MSL estimator is unbiased and efficient see the appendix.

Denote the simulated counterpart of (3.22) by $\Pr^{sim}(o_i^*, w_i^{obs})$ for simplicity, and let $\theta = (\eta, \mu, \Sigma)$ be the set of our parameters, we can write down the simulated log-likelihood

	(1)
	log hourly earnings
1-Digit SOC 2000=1	0
	(.)
1-Digit SOC 2000=2	-1.787**
	(0.667)
1-Digit SOC 2000=3	-0.238
	(0.488)
1-Digit SOC 2000=4	-1.928
	(0.985)
1-Digit SOC 2000=5	-0.274
	(0.996)
1-Digit SOC 2000=6	-0.980*
	(0.468)
1-Digit SOC 2000=7	-1.191**
	(0.435)
1-Digit SOC 2000=8	-1.589***
	(0.456)
1-Digit SOC 2000=9	-1.328**
	(0.436)
1-Digit SOC 2000=1XMathematical Skill	-1.113*
	(0.543)
1-Digit SOC 2000=2XMathematical Skill	0.801
	(0.653)
1-Digit SOC 2000=3XMathematical Skill	-0.716
	(0.369)
1-Digit SOC 2000=4XMathematical Skill	-0.138
	(0.342)
1-Digit SOC 2000=5XMathematical Skill	0.127
	(0.274)
1-Digit SOC 2000=6XMathematical Skill	-0.219
	(0.208)
1-Digit SOC 2000=7XMathematical Skill	-0.117
	(0.222)
1-Digit SOC 2000=8XMathematical Skill	0.546
	(0.278)
1-Digit SOC 2000=9XMathematical Skill	0.0841
	(0.206)
1-Digit SOC 2000=1XVerbal Skill	-0.218
	(0.267)
1-Digit SOC 2000=2XVerbal Skill	0.464**
	(0.165)
1-Digit SOC 2000=3XVerbal Skill	-0.158
	(0.198)
1-Digit SOC 2000=4XVerbal Skill	1.059
	(0.999)
1-Digit SOC 2000=5XVerbal Skill	-1.122
	(1.083)
1-Digit SOC 2000=6XVerbal Skill	-0.0785
	(0.246)
1-Digit SOC 2000=7XVerbal Skill	0.0634
	(0.249)
1-Digit SOC 2000=8XVerbal Skill	0.0653
	(0.172)
1-Digit SOC 2000=9XVerbal Skill	-0.0145
	(0.138)
male	0
	(.)
female	-0.0116
	(0.0273)
Age	0.0162*
	(0.00631)
Constant	2.739***
	(0.433)
Year Fixed Effects	Yes
Observations	265

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C1: Veta Regression

Parameter	Description	Number of Parameters	Type
μ_{kmt}	Location parameter of the subject-period-specific skill distribution.	30	Estimated
σ_{kmt}	Scale parameter of the subject-period-specific skill distribution.	30	Estimated
η_{ot}	Occupation-period specific fixed effect.	24	Estimated
ω_{ot}	Occupation-period specific occupation preferences	24	Estimated
τ	Year fixed effects	23	Estimated
γ_t	Period-specific gender coefficient	3	Estimated
ε_t	Period-specific linear experience coefficient	3	Estimated
ϕ	Standard deviation of log wage measurement errors	1	Calibrated

Table C2: Summary of Model Parameters

function of the as:

$$ll^{sim}(\theta; \phi^2) = \frac{1}{N} \sum_i \sum_{o=1}^{O} \mathbf{1}_{(o=o^*)} \ln \overset{sim}{\Pr}(o_i^*, w_i^{obs}) \tag{3.26}$$

and we can estimate θ as:

$$\hat{\theta} = \arg\max_{\theta} ll^{sim}(\theta; \phi^2) \tag{3.27}$$

So to specify the complete algorithm:

1. Set $q = 1$ and make a guess for $\hat{\theta}_1$. Specify a tolerance criterion ϵ. Set R, the number of draws used to approximate the integral to a reasonably high number.

2. For each individual i, given $\hat{\theta}_q$ draw a vector of s_i, R times, denoting each as s_i^r.

3. For $r = 1$ to R :

   (a) Calculate $\nu_i^r = w_i^{obs} - \left[\eta_{o^*} + \sum_K \lambda_{o^*k} s_{ik}^r\right]$.

   (b) For a given pair s_i^r, ν_i^r calculate $\Pr^r(o_i^*, w_i^{obs})$.

4. Average over all R values of $\Pr^r(o_i^*, w_i^{obs})$ to obtain:
   $\Pr^{sim}(o_i^*, w_i^{obs}) = \frac{1}{R} \sum_{r=1}^{R} \Pr^r(o_i^*, w_i^{obs})$.

5. Repeat steps $2 - 4$ for all N individuals. Calculate the log-likelihood via (3.26) denoting it as ll_q^{sim}.

6. If $|ll_q^{sim} - ll_{q-1}^{sim}| < \epsilon$, terminate here. Otherwise increment q and find a new value $\hat{\theta}_q$ and repeat from step 2.

Finally, for the numerical evaluation of the integral I use a grid of $10,000$ quasi-random Halton draws, which have been shown to provide about an order of magnitude more accuracy than simple random draws (Train (2009)). To ensure stochastic equicontinuity I use the same set of points for each agent at each iteration. The likelihood function generated by this problem is smooth, but not globally concave which makes it difficult for gradient-based optimization routines that are prone to converge to local minima. This is a general problem for the class of discrete choice models, but especially here given the high dimensionality of the parameter space. To maximize the log-likelihood function, I therefore use a two-step procedure: 1. I estimate the model under the assumption of no skill heterogeneity within each subject ($\sigma^2_{kmt} = 0, \forall k, m, t$). This avoids the evaluation of the integral saving considerable computational time. Taking advantage of this speed gain, I start the optimization using $1,000,000$ random starting values. 2. I take the best of these runs as a starting value to fit the full model. Optimization is performed using Matlab's `fminunc` routine. All critical values for convergence are set to $1e-6$.

11.3 Standard errors

I calculate numerical standard errors following the relationship between the hessian of the likelihood function and the information identity (c.f. Train (2009)). For the correctly specified model, the error of the MLE estimate $\hat{\theta}$ is distributed according to:

$$\sqrt{N}(\hat{\theta} - \theta^*) \to N(0, -\mathbf{H}^{-1})$$

where θ^* is the true parameter vector, and $-\mathbf{H}$ is the information matrix. To avoid complications due to the numerical procedure and the high dimensionality of the problem, I calculate a numerical hessian of the likelihood function at the SMLE estimate and then use a pseudo-inverse (c.f. Gill & King (2004)) to obtain the standard errors for the estimated parameters.

11.4 Asymptotic equivalence of SML and ML

The asymptotic properties of the simulated maximum likelihood estimator have been well understood (c.f. Gourieroux and Monfort (1993), Lee (1995), and Hajivassiliou & Ruud (1994)). This short exposition here is based on the discussion in Train (2009, Chapter 10) for simplicity. Generally maximum likelihood estimation proceeds by maximizing the log-likelihood function:

$$ll(\theta) = \sum_n \ln \Pr_n(\theta) \tag{3.28}$$

where θ is a vector of parameters and $\Pr_n(\theta)$ is the exact probability of the observed choice of observation n given θ.

Similarly, simulated maximum likelihood maximizes the simulated maximum likelihood function:

$$sll(\theta) = \sum_n \ln \overset{sim}{\underset{n}{\Pr}}(\theta) \tag{3.29}$$

where $\Pr_n^{sim}(\theta)$ is the simulated probability of the observed choice of observation n. It is known,that if $\Pr_n^{sim}(\theta)$ is an unbiased simulator for the exact probability - i.e. $E_r(\Pr_n^{sim}(\theta)) = \Pr_n(\theta)$, where the expectation is taken over r simulation draws, then there are three sources of bias in the SML estimator:

1. Sampling bias, which is the same as in the ML estimator and which goes to 0 as $N \to \infty$.
2. Simulation noise, which goes to 0 as the number of simulation draws $R \to \infty$.
3. Bias due to the fact, that $\ln \Pr_n^{sim}(\theta)$ is not an unbiased estimator of $\ln \Pr_n(\theta)$. This bias disappears if R grows faster than $\sqrt{N}$.

Hence the results, which are derived under fairly general conditions, indicate, that if $\Pr_n^{sim}(\theta)$ is an unbiased simulator, and the number of simulation draws is sufficiently larger than $\sqrt{N}$, then the MSL estimator is consistent, asymptotically normal, efficient and equivalent to traditional ML.

Therefore, the only thing that we have to show, is that our simulated joint probability $\Pr^{sim}(o_i^*, w_i^{obs}|\theta)$ is an unbiased estimator of the exact probability $\Pr(o_i^*, w_i^{obs}|\theta)$. To show this, let's remind ourselves, of how the simulated probability is obtained:

$$\overset{sim}{\Pr}(o_i^*, w_i^{obs}|\theta) = \frac{1}{R}\sum_{r=1}^{R}\left\{\left(\frac{e^{\eta_{o^*}+\sum_K \lambda_{o^*k}s_{ik}^r}}{\sum_{o=1}^{O} e^{\eta_o+\sum_K \lambda_{ok}s_{ik}^r}}\right)\left(\frac{e^{(-\frac{\left(w_i^{obs}-[\eta_{o^*}+\sum_K \lambda_{o^*k}s_{ik}^r]\right)^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}}\right)\right\} \tag{3.30}$$

where s_i^r is the rth simulation draw of s_i. Compare this to the exact probability:

$$\Pr(o_i^*, w_i^{obs}|\theta)) = \int\left\{\left(\frac{e^{\eta_{o^*}+\sum_K \lambda_{o^*k}s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o+\sum_K \lambda_{ok}s_{ik}}}\right)\left(\frac{e^{(-\frac{\left(w_i^{obs}-[\eta_{o^*}+\sum_K \lambda_{o^*k}s_{ik}]\right)^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}}\right)\right\} f(s)d(s) \tag{3.31}$$

Now

$$E\left(\overset{sim}{\Pr}(o_i^*, w_i^{obs}|\theta)\right) =$$

$$E\left[\frac{1}{R}\sum_{r=1}^{R}\left\{\left(\frac{e^{\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}^r}}{\sum_{o=1}^{O} e^{\eta_o+\sum_K \lambda_{ok}s_{ik}^r}}\right)\left(\frac{e^{(-\frac{\left(w_i^{obs}-[\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}^r]\right)^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}}\right)\right\}\right]$$

$$= \frac{1}{R}\sum_{r=1}^{R} E\left\{\left(\frac{e^{\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}^r}}{\sum_{o=1}^{O} e^{\eta_o+\sum_K \lambda_{ok}s_{ik}^r}}\right)\left(\frac{e^{(-\frac{\left(w_i^{obs}-[\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}^r]\right)^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}}\right)\right\}$$

$$= \frac{1}{R}\sum_{r=1}^{R}\int\left\{\left(\frac{e^{\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o+\sum_K \lambda_{ok}s_{ik}}}\right)\left(\frac{e^{(-\frac{\left(w_i^{obs}-[\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}]\right)^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}}\right)\right\} f(s)d(s)$$

$$= \int\left\{\left(\frac{e^{\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}}}{\sum_{o=1}^{O} e^{\eta_o+\sum_K \lambda_{ok}s_{ik}}}\right)\left(\frac{e^{(-\frac{\left(w_i^{obs}-[\eta_{o*}+\sum_K \lambda_{o*k}s_{ik}]\right)^2}{2\phi^2})}}{\sqrt{2\pi\phi^2}}\right)\right\} f(s)d(s)$$

$$= \Pr(o_i^*, w_i^{obs}|\theta)) \quad (3.32)$$

the third line follows from the definition of the expected value operator and the fact that all s_i^r are i.i.d.

Hence, we have shown, that $\Pr^{sim}(o_i^*, w_i^{obs}|\theta)$ is an unbiased estimator for $\Pr(o_i^*, w_i^{obs}|\theta)$. Furthermore, $\Pr^{sim}(o_i^*, w_i^{obs}|\theta)$ is a continuous and twice differentiable function.

11.5 Sidenote on "effective skills"

I wanted to highlight some potential limitations of the approach taken in this paper. Namely, my estimation is only ever going to recover *effective* skills. To explain what this means for our estimation strategy, consider the following simple one-dimensional example:

So far we have assumed that output was produced by the combination of the worker's skills and the occupation's task requirements:

$$\ln(Y) = \lambda_o s_i \quad (3.33)$$

Now this expression is observationally equivalent to another expression:

$$\ln(Y) = \chi\lambda_o s_i^* \tag{3.34}$$

$$s_i^* = \frac{s_i}{\chi} \tag{3.35}$$

where χ is a general productivity parameter that is common to all occupations. In a one-period case this is not particularly troublesome, but in a multi-period setting one might want to take the possibility of the general productivity of a certain skill changing seriously. Unfortunately, it is not possible here to address this issue and disentangle the two, which would require some information on the evolution of χ. Therefore, I just wanted to clarify that what the estimation procedure recovers is $s_i = s_i^*\chi$ which I dub *effective* skills.

Despite this caveat, I believe that there is not too much to be concerned about here - i.e. I think that I can justify that $\chi_{kt} \approx 1\ \forall k, t$ for the following reasons: i) We have observed that task requirements have changed over the period, which means that at least part of χ is actually observed and thus controlled for. ii) The inclusion of occupation-time-specific fixed effects will soak up some of this aggregate change.

SUMMARY AND FUTURE RESEARCH

This thesis has studied wealth and human capital inequality as well as the interactions between these two important dimensions of economic inequality. I found that skill inequality is an important driver of wealth inequality particularly when it interacts with technological change at the aggregate level. Similarly, I found that wealth inequality can be an important contributing factor to differences in human capital accumulation - an effect that is amplified when the economy experiences large unexpected shocks. Finally, I showed that the supply of skills is more varied and more changeable than we might have expected, depending ultimately on a combination of ability and talent, institutional factors and market forces, individual opportunities and a little bit of luck.

While each chapter necessarily only addressed a small part of a much larger whole, they do draw from each other and make important connections across the pages of this thesis. Chapter 1 shows that changes in income risk brought about by technological change have knock-on effects on the distribution of wealth, setting the scene for considering the endogenous human capital distribution in Chapter 2. The latter chapter adds to this the reverse direction, as it shows, that the distribution of wealth can influence the distribution of human capital and therefore income. This channel is shown to be important during large economic shocks. Chapter 3 looks closer at a thus far unexplored dimension of human capital inequality, by studying the multidimensional skill distributions of university graduates. In the same vein as the analysis in Chapter 1, it suggests that the distribution of human capital responds to technological change.

Chapter 1

In this chapter, I proposed a model to study the effect of differentiated, cognitive skill-biased technological change on income and wealth inequality in the UK. My model combined elements of the "Task-Skill" literature with a heterogeneous agent incomplete market model. The model introduced a novel income process that relates worker's earnings to their multidimensional skill set and occupation-specific returns to those skills. Utilising measures of cognitive and non-cognitive skills from a comprehensive panel dataset for the UK (Understanding Society), the calibrated model captured many features of the income process observed in data and provides additional features beyond other approximation techniques, such as the possibility for technological change to directly affect income risk. In an application, I used the model to assess the impact of cognitive skill-biased technological change due to increased computer usage in the UK over 1980-2016. My results suggested that this specific type of technological change can account for most increases in labour income inequality observed over that period and is consistent with stylized facts about changes to wealth inequality over the same period.

Chapter 2

In this chapter, I developed a general equilibrium heterogeneous agent incomplete market model with endogenous wealth and human capital to analyze the interactions between these two factors. The model was calibrated to the UK economy in the pre-Covid-19 period using Understanding Society data. In terms of methods, this chapter extended the standard incomplete markets model by allowing workers to invest in "risky" human capital resulting in an endogenously generated joint distribution of wealth, human capital and income. This constitutes a novel contribution as it allows for analysis of how initial differences in household wealth lead to differences in human capital accumulation, by otherwise identical households.

My main findings are that there exist important non-linearities in human capital investments, with workers with low levels of wealth investing considerably less in accumulating human capital than their counterparts with more wealth. I further analyzed the economic dynamics of the distribution of human capital in the aftermath of an unexpected economic shock, showing that wealth-poorer households are more exposed to these shocks, implying that the distribution of wealth matters for the recovery of the economy following recessions. This is a novel finding, as generally the distribution of wealth is seen as not to matter much for the dynamics of the aggregate economy.

The chapter further assesses the impact of the Covid-19 pandemic and associated support measures in the UK. The model predicts that the UK economy will likely suffer a significant reduction in human capital in the aftermath of the Covid-19 pandemic but targeted policy action has helped reduce the impact, particularly for low-wealth and otherwise vulnerable households.

Chapter 3

In this chapter, I studied the distribution of skills among university graduates in the UK. For this purpose, I developed a new model of occupational choice and wage determination for university graduates in the UK. The analysis focuses on two types of general skills: mathematical/technical and verbal/organizational. The model was structurally estimated using UK Labour Force data covering the period from 1994 to 2019 and I find quantitative evidence of changing multivariate skill distributions over time. In terms of methods, this chapter is novel in its use of a structural model to estimate graduates' skills and thus offers an innovative way of thinking about the return to different university subjects. It further presents a novel set of stylized facts that provide some insight as to the qualitative dimension of the skill supply at the top end of the distribution.

My main findings are that between 1994 and 2019, the typical graduate's level of mathematical skills increased considerably, while verbal skills decreased in turn. This trend is driven by increasing specialization for STEM and Business & Economics degrees and increasing generalization among Arts & Humanities and Other Subjects. For most graduates, mathematical/technical skills have become the single biggest contributing factor to their earnings, making up around 50% of their hourly wage compared to 27% in the mid-90s. The results suggest that graduate skill supply has adjusted to changing labour market requirements.

Future Research

This thesis has investigated three main mechanisms that play a role in defining economic inequality in the past and going forward: i) technological change affects wealth inequality (Chapter 1); ii) wealth and human capital inequality affect each other (Chapter 2); and iii) technological change affects the distribution of human capital (Chapter 3). Each of these mechanisms is important in and by itself and probably deserves a much deeper treatment than this thesis could provide. But there are a number of potential avenues for further research that suggest themselves based on the analysis of the preceding pages:

One promising extension would be to apply the model developed in Chapter 2 to a study of technological change. The seeds for this were sown in Chapter 1, and it would be very interesting to see how technological change has affected the joint distribution of wealth and human capital. Such an analysis could then be extended into the future to make predictions about the impact of future technologies.

Another important area of application for the model would be policy analysis. Since the model features endogenous wealth and human capital it can answer important counterfactual questions related to diverse policy issues that are usually not studied together. For example, it can assess the impact of a wealth tax on human capital inequality or the effect of a training subsidy on wealth inequality. A further application to study the optimal taxation schedule to repay the Covid-19 debt has already been suggested in the text.

Finally, the model in Chapter 3 does not feature wealth, but clearly going to university and choosing a specific subject is a risky decision with uncertain outcomes and thus wealth should play a role in influencing this decision. Thus introducing assets into the model and evaluating the role that differences in wealth play at this crucial moment in a young person's life would be a worthwhile extension.

I do leave these and all other suggestions for future research.